ATM Switch
Configuration Manual

CABLETRON SYSTEMS

The Complete Networking Solution™

Notice

Cabletron Systems reserves the right to make changes in specifications and other information contained in this document without prior notice. The reader should in all cases consult Cabletron Systems to determine whether any such changes have been made.

The hardware, firmware, or software described in this manual is subject to change without notice.

IN NO EVENT SHALL CABLETRON SYSTEMS BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS MANUAL OR THE INFORMATION CONTAINED IN IT, EVEN IF CABLETRON SYSTEMS HAS BEEN ADVISED OF, KNOWN, OR SHOULD HAVE KNOWN, THE POSSIBILITY OF SUCH DAMAGES.

© Copyright April 1996 by:

Cabletron Systems, Inc.
35 Industrial Way
Rochester, NH 03867-5005

All Rights Reserved
Printed in the United States of America

Order Number: 9031917

SPECTRUM, **Remote LANVIEW**, and **LANVIEW** are registered trademarks and **MMAC-Plus** is a trademark of Cabletron Systems, Inc.

Ethernet is a trademark of Xerox Corporation.

FCC Notice

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment uses, generates, and can radiate radio frequency energy and if not installed in accordance with the operator's manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause interference in which case the user will be required to correct the interference at his own expense.

WARNING: Changes or modifications made to this device which are not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.

VCCI Notice

This equipment is in the 1st Class Category (information equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council for Interference by Information Technology Equipment (VCCI) aimed at preventing radio interference in commercial and/or industrial areas.

Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, etc.

Read the instructions for correct handling.

この装置は、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（VCCI）基準に適合しております。
従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。
取扱説明書に従って正しい取り扱いをして下さい。

DOC Notice

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le présent appareil numérique n'émet pas de bruits radioélectriques dépassant les limites applicables aux appareils numériques de la class A prescrites dans le Règlement sur le brouillage radioélectrique édicté par le ministère des Communications du Canada.

Table of Contents

Preface

Chapter Summaries xiii
Technical Support xiv
Typographical Styles xv
Important Information Indicators xv
Laser Warning xvii
Safety Agency Compliance xviii
Safety Precautions xviii
Symbols xviii
Modifications to Equipment xix
Placement of a FORE Systems Product xix
Power Cord Connection xix

Chapter 1 Configuring PVCs

1.2 Virtual Paths 1-3
1.2.1 Through Paths 1-4
1.2.2 Originating and Terminating Paths 1-6
1.3 Listing Virtual Paths 1-7
1.4 Virtual Channels 1-10
1.4.1 Listing Virtual Channels 1-13
1.5 Creating Permanent Virtual Circuits (PVCs) 1-15
1.5.1 Creating a Virtual Path 1-15
1.5.1.1 Terminating a PVC at a Switch 1-18
1.5.1.2 Creating ATM ARP Entries 1-19
1.5.1.3 Listing ATM ARP Entries 1-20
1.5.2 Creating a Virtual Channel 1-21
1.6 Traffic Types 1-23
1.7 Traffic Policing (Usage Parameter Control) 1-24
1.7.1 Setting the CLP Bit 1-24
1.7.2 Leaky Bucket Algorithm 1-24
1.7.3 UNI 3.0 UPC Traffic Contract Parameters 1-25
1.7.4 AMI UPC Commands 1-27

Chapter 2 Configuring a Classical IP ATM Network

2.1.1 Logical IP Subnets 2-2
2.1.2 Classical IP Interfaces 2-2
2.1.3 SPANS Interface 2-3
2.2 Address Registration and ILMI 2-4
2.2.1 NSAP Addresses 2-4
2.2.2 Operating with ILMI Support 2-5
2.2.3 Operating without ILMI Support 2-5
2.2.4 Configuration 2-5
2.3 ARP and ARP Servers 2-6

2.3.1 Theory 2-6
2.3.2 Configuring a Host to be an ARP Server 2-7
2.3.2.1 Configuring a Host as an ARP Server for One-Time Use 2-8
2.3.2.2 Configuring a Host as a Permanent ARP Server 2-9
2.3.2.3 Configuring a Hewlett-Packard Computer as an ARP Server 2-10
2.3.3 Configuring a Cabletron Switch to be an ARP Server 2-11
2.3.4 Classical IP Operation 2-12
2.3.5 Operational Issues 2-13
2.4 Classical IP PVCs 2-14
2.4.1 Theory and Configuration 2-14
2.4.2 Revalidation and Removal 2-15
2.5 Debugging 2-16
2.6 Configuring the Network 2-17
2.6.1 Third-Party Host with No ILMI and No RFC-1577 Support 2-18
2.6.2 Third-Party Switch with ILMI and No RFC-1577 Support 2-19
2.6.3 Third-Party Switch with RFC-1577 and No ILMI Support 2-20

Chapter 3 Configuring an Emulated LAN

3.1 Introduction 3-1
3.2 Emulated LAN Components 3-2
3.2.1 LAN Emulation Client (LEC) 3-3
3.2.2 LAN Emulation Configuration Server (LECS) 3-3
3.2.3 LAN Emulation Server (LES) 3-3
3.2.4 Broadcast and Unknown Server (BUS) 3-3
3.3 Emulated LAN Operation 3-4
3.3.1 Initialization 3-6
3.3.2 Registration and Address Resolution 3-6
3.3.3 Data Transfer 3-7
3.4 Configuring an ELAN 3-8
3.4.1 Configuring an LECS Configuration Database File 3-9
3.4.1.1 Before You Begin 3-9
3.4.1.2 LECS Configuration File Syntax 3-10
3.4.1.3 Defining an ELAN 3-12
3.4.1.4 Defining a Client 3-14
3.4.1.5 LECS Control Parameters 3-15
3.4.2 Sample LECS Configuration File 3-16
3.4.3 The Default LECS Configuration File 3-19
3.4.4 Starting the LAN Emulation Services 3-21
3.4.4.1 Starting the LECS 3-21
3.4.4.2 Creating a LES and a BUS 3-23
3.4.5 Starting the LEC(s) and Joining an ELAN 3-24
3.4.5.1 Creating a LEC 3-24
3.4.5.2 Configuring the LEC Failover Mechanism 3-26

Chapter 4 SONET Configuration

4.1.1 Transmit Indicators 4-1
4.1.2 Receive Indicators 4-2
4.2 Configuring SONET Mode 4-2
4.3 Configuring SONET Empty Cells 4-3
4.4 Configuring SONET Loopback 4-4
4.4.1 Diagnostic Loopback 4-4
4.4.2 Line Loopback 4-4
4.5 Displaying SONET Error Counters 4-6
4.6 SONET Error Counter Descriptions 4-7
4.7 Configuring SONET Timing 4-10

Chapter 5 DS-3 Configuration

5.1.1 Transmit Indicators 5-1
5.1.2 Receive Indicators 5-1
5.2 Configuring DS-3 Mode 5-2
5.3 Configuring DS-3 Empty Cells 5-3
5.4 Configuring DS-3 Line Length 5-3
5.5 Configuring DS-3 Framing 5-4
5.6 Configuring DS-3 Payload Scrambling 5-5
5.7 Configuring DS-3 Loopback 5-6
5.7.1 Cell Loopback 5-7
5.7.2 Payload Loopback 5-7
5.7.3 Diagnostic Loopback 5-7
5.7.4 Line Loopback 5-7
5.8 Displaying DS-3 Error Counters 5-8
5.9 DS-3 Error Counter Descriptions 5-9
5.10 Configuring DS-3 Timing 5-11

Chapter 6 E-3 Configuration

6.1.1 Transmit Indicators 6-1
6.1.2 Receive Indicators 6-1
6.2 Configuring E-3 Mode 6-2
6.3 Configuring E-3 Payload Scrambling 6-2
6.4 Configuring E-3 Loopback 6-3
6.4.1 Cell Loopback 6-4
6.4.2 Payload Loopback 6-4
6.4.3 Diagnostic Loopback 6-4
6.4.4 Line Loopback 6-4
6.5 Configuring E-3 Empty Cells 6-5
6.6 Displaying E-3 Error Counters 6-6
6.7 E-3 Error Counter Descriptions 6-7
6.8 Configuring E-3 Timing 6-9

Chapter 7 TP25 Configuration

8.1.1 Transmit Indicators 7-1
8.1.2 Receive Indicators 7-1
8.2 Configuring TP25 Loopback 7-2
8.2.1 Line Loopback 7-2
8.3 Displaying TP25 Error Counters 7-4
8.4 TP25 Error Counter Descriptions 7-4

Chapter 8 TAXI Configuration

9.1.1 Transmit Indicators 8-1
9.1.2 Receive Indicators 8-1
9.2 Configuring TAXI Loopback 8-2

Appendix A AMI Overview

A.1.1 Login from Serial Port A-3
A.1.2 Login from Telnet A-4
A.2 AMI Commands Not Available When Running Remotely A-6
A.3 AMI Root Menu for an Open Session A-7
A.3.1 About Command A-8
A.3.2 Close Command A-9
A.3.3 Configuration Commands A-10
A.3.4 Exit Command A-10
A.3.5 Help Command A-10
A.3.6 History Command A-11
A.3.7 Open Command A-12
A.3.8 Operation Commands A-13
A.3.9 Ping Command A-13
A.3.10 Redo Command A-14
A.3.11 Rows Command A-16
A.3.12 Statistics Commands A-16
A.3.13 Top Command A-16
A.3.14 Up Command A-16

Appendix B AMI Configuration Commands

B.1.1 Displaying Alarm Conditions B-2
B.1.2 Enabling an Alarm B-3
B.1.3 Disabling an Alarm B-4
B.1.4 Resetting an Alarm B-5
B.2 ATM ARP Configuration Commands B-6
B.2.1 ARP Server Configuration Commands B-6
B.2.2 Deleting an ARP Entry B-7
B.2.3 Flushing the ATM ARP Cache B-8
B.2.4 Getting the NSAP Address for a Classical IP Interface B-8

B.2.5 Creating an IP to NSAP Address Mapping B-8
B.2.6 Creating a Classical IP PVC B-9
B.2.7 Creating a FORE IP PVC ARP Entry B-9
B.2.8 Displaying the ATM ARP Entries B-10
B.3 Switch Board Configuration Commands B-11
B.3.1 Configuring the Clock Scaling Factor on a Switch Board B-11
B.3.2 Displaying the Board Configuration B-12
B.3.3 Displaying the Board Topology B-13
B.4 IP Configuration Commands B-14
B.4.1 Configuring the IP Address B-14
B.4.2 Configuring the IP State B-15
B.4.3 Configuring the IP Broadcast Address B-15
B.4.4 Configuring IP Forwarding B-16
B.4.5 Configuring the IP Subnet Mask B-16
B.4.6 Configuring IP Routes B-17
B.4.7 Displaying the IP Interface Configuration B-19
B.5 LAN Emulation Configuration Commands B-20
B.5.1 Broadcast and Unknown Server (BUS) Configuration Commands B-20
B.5.2 LAN Emulation Client (LEC) Configuration Commands B-24
B.5.3 LAN Emulation Configuration Server (LECS) Commands B-31
B.5.4 LAN Emulation Server (LES) Configuration Commands B-36
B.6 Network Module Configuration Commands B-42
B.6.1 Displaying Network Module Configuration Information B-42
B.6.2 Configuring Distributed Timing on a Network Module B-43
B.6.3 Configuring Traffic on a Network Module B-48
B.7 NSAP Configuration Commands B-53
B.7.1 NSAP Route Configuration Commands B-53
B.7.2 NSAP Prefix Configuration Commands B-56
B.7.3 NSAP ILMI Configuration Command B-59
B.8 Port Configuration Commands B-60
B.8.1 CDVT Port Configuration Commands B-60
B.8.2 E-3 Port Configuration Commands B-61
B.8.3 TP25 Port Configuration Commands B-66
B.8.4 DS-3 Port Configuration Commands B-68
B.8.5 J-2 Port Configuration Commands B-73
B.8.6 Port Policing Configuration Command B-76
B.8.7 SONET Port Configuration Commands B-77
B.8.8 Showing the Port Configuration B-82
B.8.9 TAXI Port Configuration Commands B-85
B.8.10 Traffic Port Configuration Commands B-87
B.8.11 VBROB Port Configuration Commands B-93
B.8.12 VBRBuffOB Port Configuration Commands B-93
B.9 Serial Port Configuration Commands B-94
B.9.1 Displaying Serial Port Information B-94
B.10 SNMP Configuration Commands B-95
B.10.1 Configuring the SNMP Community Access B-95
B.10.2 Configuring SNMP Traps B-95
B.11 SPANS Configuration Commands B-98
B.11.1 Deleting a SPANS Signalling Path B-98
B.11.2 Creating a SPANS Signalling Path B-99

B.11.3 Showing the SPANS Signalling Path Configuration....B-101
B.12 SPVC Configuration Commands....B-104
B.12.1 Deleting an SPVC....B-104
B.12.2 Creating an SPVC....B-105
B.12.3 Displaying SPVC Information....B-107
B.13 Switch Configuration Commands....B-108
B.13.1 Setting or Changing the Switch Name....B-108
B.13.2 Setting the Minimum Number of Reserved VCIs for PMPs....B-109
B.13.3 Setting the Maximum Number of Reserved VCIs for PMPs....B-110
B.13.4 Displaying the Switch Configuration....B-111
B.14 System Configuration Commands....B-112
B.14.1 Displaying System Information....B-112
B.14.2 System Log Configuration Commands....B-113
B.14.3 AMI Timeout Configuration Command....B-117
B.14.4 Configuring the Units for UPC Contracts....B-117
B.15 Topology Configuration Commands....B-118
B.15.1 ForeThought PNNI Configuration Commands....B-118
B.15.2 SPANS Topology Configuration Commands....B-126
B.16 UNI 3.0 Configuration Commands....B-128
B.16.1 Deleting a UNI 3.0 Signalling Path....B-128
B.16.2 Creating a UNI 3.0 Signalling Path....B-129
B.16.3 Displaying UNI 3.0 Signalling Paths....B-131
B.17 Usage Parameter Control Configuration Commands....B-134
B.17.1 Deleting a UPC Traffic Contract....B-134
B.17.2 Creating a UPC Traffic Contract....B-135
B.17.3 Displaying the UPC Traffic Contracts....B-137
B.18 Virtual Channel Configuration Commands....B-138
B.18.1 Deleting a Virtual Channel....B-138
B.18.2 Modifying a Virtual Channel....B-139
B.18.3 Creating a Virtual Channel....B-139
B.18.4 Displaying the Virtual Channel Configuration....B-141
B.19 Virtual Path Configuration Commands....B-143
B.19.1 Deleting a Virtual Path....B-143
B.19.2 Modifying a Virtual Path....B-144
B.19.3 Creating a Virtual Path....B-145
B.19.4 Displaying Virtual Paths....B-149

Appendix C AMI Operation Commands

C.1.1 Backing Up the Database....C-2
C.1.2 Initializing the Database....C-3
C.1.3 Resetting the Database....C-4
C.1.4 Restoring the Database....C-4
C.2 Environment Commands....C-5
C.2.1 CPU Operation....C-5
C.2.2 Switch Fabric Operation....C-6
C.2.3 Showing Switch Fabric Temperature Information....C-6
C.2.4 Configuring the Switch Fabric Temperature Thresholds....C-7
C.2.5 Fan Operation....C-8

C.2.6 Power Supply Operation C-9
C.2.7 Temperature Sensor Operation C-10
C.2.8 Panic Acknowledgment Commands C-11
C.2.9 Clearing the Panic Flag C-11
C.2.10 Displaying the Panic Dump File C-12
C.3 Displaying and Setting the Date and Time C-14
C.4 FLASH Operation Commands C-15
C.4.1 Copying a File to FLASH Memory C-15
C.4.2 Deleting a File from FLASH Memory C-16
C.4.3 Displaying the FLASH Memory Directory C-17
C.4.4 Displaying Free Space on the FLASH File C-17
C.4.5 Getting a FLASH File C-17
C.4.6 Initializing the FLASH File C-18
C.4.7 Putting a FLASH File on a Remote Host C-18
C.4.8 Renaming a FLASH File C-19
C.5 Setting or Changing the Password C-19
C.6 Upgrading the Switch C-20
C.7 Displaying and Changing the Version of Software C-21
C.8 Rebooting the Switch C-21

Appendix D AMI Statistics Commands

D.2 AAL4 Statistics D-2
D.3 AAL5 Statistics D-3
D.4 ATM Statistics D-4
D.5 Switch Board Statistics D-5
D.6 Control Port Statistics D-5
D.7 DS-3 Statistics D-6
D.8 E-3 Statistics D-9
D.9 ICMP Statistics D-12
D.10 Interface Statistics D-15
D.11 IP Statistics D-17
D.12 Network Module Statistics D-20
D.13 Port Statistics D-22
D.14 SONET Statistics D-23
D.15 SPANS Statistics D-26
D.16 TCP Statistics D-27
D.17 TP25 Statistics D-29
D.18 UDP Statistics D-30
D.19 UNI 3.0 Statistics D-31
D.20 VCC Statistics D-32
D.21 VPC Statistics D-33

Appendix E SNMP Configuration

E.2 SNMP Traps E-3
E.2.1 Adding SNMP Trap Destinations E-11
E.2.2 Displaying SNMP Trap Destinations E-11
E.2.3 Removing SNMP Trap Destinations E-12

Appendix F ForeThought PNNI

F.1.1 Hello Protocol F-2
F.1.2 Topology Database Exchange F-2
F.1.3 Flooding F-2
F.1.4 Hierarchical Routing F-3
F.2 The Physical Network F-5
F.2.1 Peer Groups F-8
F.2.2 Peer Group Topology F-8
F.2.3 Border Switches F-8
F.2.4 Peer Group Summary Node (PGSN) F-9
F.2.5 Backbone Topology F-9
F.2.6 Single Switch Perspective F-9

PREFACE

This manual provides the technical information needed to configure the ATM Switches, the LAN and WAN network modules, and the accompanying software. This document also provides general ATM information and general product information. This document was created for users with various levels of experience. If you have any questions or problems, please contact Cabletron Systems' Technical Support.

Chapter Summaries

Chapter 1 - Configuring PVCs - Describes how to create PVCs on a switch through the ATM Management Interface (AMI).

Chapter 2 - Configuring a Classical IP ATM Network - Describes how to design, configure, and maintain a Classical IP ATM network.

Chapter 3 - Configuring an Emulated LAN - Provides an overview of LAN Emulation and gives an example of how to configure an Emulated LAN.

Chapter 4 - SONET Configuration - Contains configuration information for supporting SONET network modules.

Chapter 5 - DS-3 and DS-1 Configuration - Contains configuration information for supporting DS-3 and DS-1 network modules.

Chapter 6 - E-3 and E-1 Configuration - Contains configuration information for supporting E-3 and E-1 network modules.

Chapter 7 - J-2 Configuration - Contains configuration information for supporting J-2 network modules.

Chapter 8 - TP25 Configuration - Contains configuration information for supporting TP25 network modules.

Chapter 9 - TAXI Configuration - Contains configuration information for supporting TAXI network modules.

Appendix A - AMI Overview - Provides an overview of AMI and contains a text and graphical description of the root level AMI commands.

Appendix B - AMI Configuration Commands - Contains a text and graphical description of the configuration level AMI commands.

Appendix C - AMI Operation Commands - Contains a text and graphical description of the operation level AMI commands.

Appendix D - AMI Statistics Commands - Contains a text and graphical description of the statistics level AMI commands.

Appendix E - SNMP Configuration - Describes the remote SNMP configuration of a switch.

Appendix F - PNNI - Describes how this scalable routing and signalling protocol can be used to simplify large network topologies.

Technical Support

If you need additional support with these products, or if you have any questions, comments or suggestions concerning this manual, feel free to contact Cabletron Systems Technical Support:

By phone:	(603) 332-9400
By CompuServe:	GO CTRON from any ! prompt
By Internet mail:	support@ctron.com
By mail:	Cabletron Systems, Inc. P.O. Box 5005 Rochester, NH 03867-0505

Typographical Styles

Throughout this manual, all specific commands meant to be entered by the user will appear on a separate line in bold typeface. In addition, use of the Enter or Return key will be represented as <ENTER>. The following example demonstrates this convention.

cd /usr <ENTER>

Commands or file names that appear within the text of this manual will be represented in the following style: "...the fore_install program will install this distribution".

As in the following example, any messages that appear on your screen during software installation and network interface administration will appear in Courier font to distinguish them from the rest of the text.

```
.... Are all four conditions true?
```

Important Information Indicators

To call your attention to safety and otherwise important information that must be reviewed to ensure correct and complete installation, as well as to avoid damage to the ATM Switch or your system, this manual utilizes the following *WARNING*/*CAUTION*/*NOTE* indicators.

WARNING statements contain information that is critical to the safety of the operator and/or the system. Do not proceed beyond a ***WARNING*** statement until the indicated conditions are fully understood or met. This information could prevent serious injury to the operator, damage to the Switch, the system, or currently loaded software, and will be indicated as follows:

WARNING!

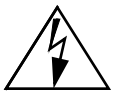

Hazardous voltages are present. To reduce the risk of electrical shock and danger to personal health, follow the instructions carefully.

Information contained in **CAUTION** statements is important for proper

installation/operation. **CAUTION** statements can prevent possible equipment damage and/or loss of data and will be indicated as:

CAUTION

You risk damaging your equipment and/or software if you do not follow these instructions.

Information contained in **NOTE** statements has been found important enough to be called to the special attention of the operator and will be set off from the text as follows:

NOTE: It is strongly recommended that you disconnect the serial cable once you have configured the ATM switch and then access the switch over the ATM network.

Laser Warning

Class 1 Laser Product:
This product conforms to applicable requirements of 21 CFR 1040 at the date of manufacture.

Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation.

The following network modules contain Class 1 lasers:

NM-2/155SMSRA-1	NM-2/155SMSRB-1	NM-2OC3/SMSRA-1
NM-4/155SMSRA-1	NM-4/155SMSRB-1	NM-4OC3/SMSRA-1
NM-4OC3/SMSRA-1A	NM-4OC3/SMSRB-1	NM-2/155SMSRC
NM-4/155SMSRC	NM-2/155SMLRA-1	NM-2/155SMLRB-1
NM-2/155SMLRC-1	NM-2OC3/SMLRB-1	NM-2OC3/SMMRA-1
NM-4/155SMLRB-1	NM-4/155SMLRC-1	NM-4OC3/SMLRB-1
NM-4OC3/SMMRA-1	NM-1/622SMIRC	NM-4/155LR3SCC
NM-4/155SR3SCC		

Safety Agency Compliance

This preface provides safety precautions to follow when installing a product.

Safety Precautions

For your protection, observe the following safety precautions when setting up your equipment:

- Follow all warnings and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source matches the voltage and frequency inscribed on the equipment's electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous voltages may be present. Conductive foreign objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

Symbols

The following symbols appear in this book.

CAUTION If instructions are not followed, there is a risk of damage to the equipment.

WARNING! Hazardous voltages are present. If the instructions are not heeded, there is a risk of electrical shock and danger to personal health.

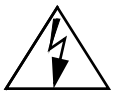

Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. Cabletron Systems, Inc., is not responsible for regulatory compliance of a modified product.

Placement of a Product

CAUTION

To ensure reliable operation of the product and to protect it from overheating, openings in the equipment must not be blocked or covered. The product should never be placed near a radiator or heat register.

Power Cord Connection

WARNING!

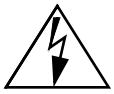

These products are designed to work with single-phase power systems having a grounded neutral conductor. To reduce the risk of electrical shock, do not plug these products into any other type of power system. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.

WARNING!

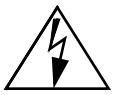

This product is shipped with a grounding type (3-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.

CHAPTER 1 Configuring PVCs

In order to interoperate with other vendors' equipment, it is often necessary to create Permanent Virtual Circuits (PVCs) on a Cabletron switch. This can be accomplished through the ATM Management Interface (AMI). This chapter will discuss the creation of PVCs through AMI on a Cabletron switch.

1.1 General Concepts

Each ATM cell contains a virtual path identifier (VPI) and a virtual channel identifier (VCI) as part of its five-byte ATM header. The VPI and VCI are used to route the cell through the ATM network. When a switch fabric receives a cell, it examines the ATM header to determine the correct output port, VPI, and VCI for the cell. For example, an ATM switch fabric can be configured such that any cell received on port A1 with VPI | VCI = 0 | 32 is switched to port B2 with VPI | VCI = 0 | 35. The translation from input port, VPI, and VCI to output port, VPI, and VCI is achieved via a mapping table in the switch fabric's memory.

The VCI value of cells does not change as the cell is switched through the ATM network via a virtual path. In a single switch environment, a cell's VPI and VCI are translated only once, but in a multiple switch environment a cell's VPI and VCI are translated many times. It is important to remember that a cell's VPI and VCI are of local significance only (i.e., link-by-link). It is also important to note that virtual connections are unidirectional; that is, they are valid in one direction only. The VPI and VCI **may** change as the cell is switched through the network.

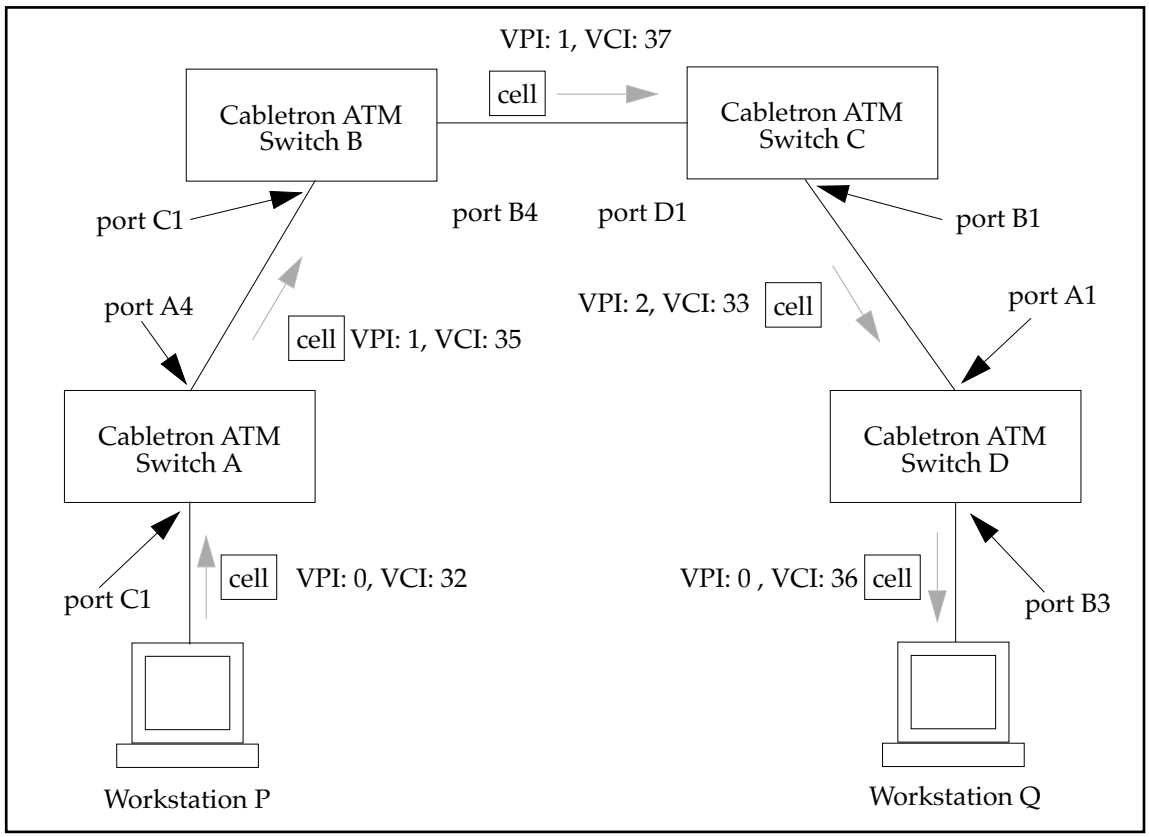

Figure 1.1 - The Cell

The set of mappings in the ATM network used to route cells from a source to a destination are generally referred to as virtual channels and virtual paths. The switch control software will normally use the SPANS protocol, FORE Systems' pre-standard signalling protocol, to create the virtual channels and virtual paths needed to establish a connection between two ATM endpoints. Not all ATM equipment supports the SPANS protocol, so it is sometimes necessary to create virtual paths and virtual channels manually in order to establish a connection between two ATM endpoints. The focus of the following sections is to explain how to create the necessary mappings to establish virtual paths and virtual channels in a network of Cabletron switches.

1.2 Virtual Paths

Virtual paths are used to establish connections between two switch fabrics in an ATM network. Once the switch fabrics are connected via a virtual path, they can use this virtual path to route virtual channels.

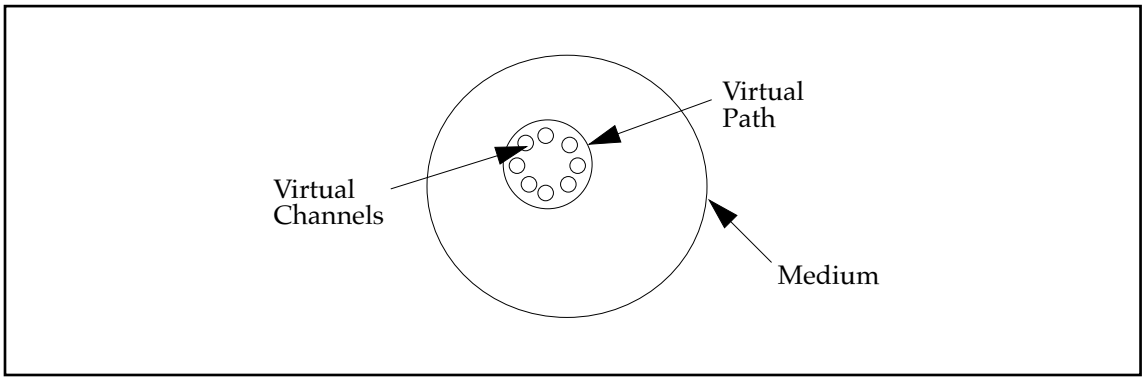

Figure 1.2 - Virtual Channels in a Virtual Path

A single virtual path can be used to route many virtual channels through the ATM network. Because a virtual path simply routes virtual channels through the network, a cell is guaranteed to have the same VCI when it exits the virtual path as it had when it entered the virtual path.

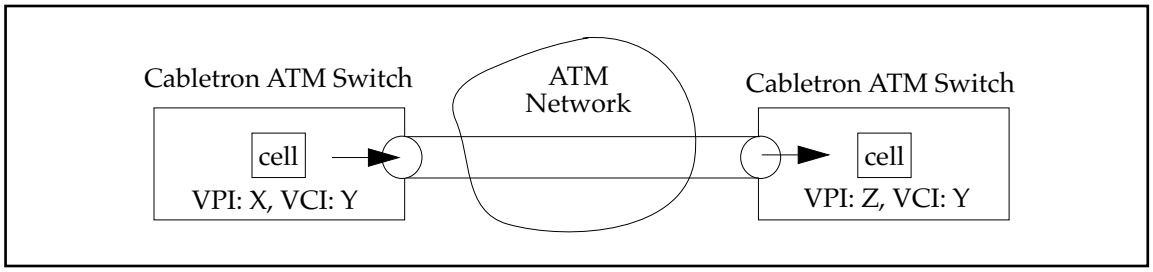

Figure 1.3 - An Example of a Virtual Path

The VCI value of cells does not change as the cell is switched through the ATM network via a virtual path. Each virtual path must originate at a switch fabric, pass through zero or more switch fabrics and terminate at another switch fabric. The origination and termination points are referred to as originating and terminating paths. Virtual paths are switched through switch fabrics via through paths. Virtual paths are made up of an originating path, zero or more through paths, and a terminating path.

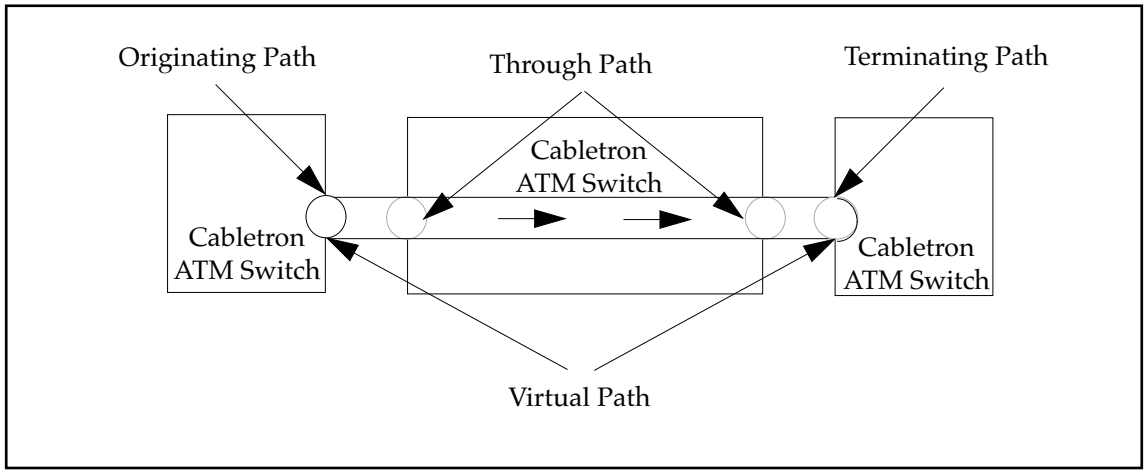

Figure 1.4 - Composition of a Virtual Path

1.2.1 Through Paths

Through paths route an entire virtual path through an ATM switch fabric. When a cell is received by a switch fabric on a through path, the VPI is examined to determine the output port and VPI. The VCI component of the ATM header remains unchanged and can have any value. So, all of the channels within the through path will be switched correctly without altering the VCI value of cells on these channels.

Four parameters are needed to define a through path on a Cabletron switch fabric: input port, input VPI, output port, and output VPI. Through paths are represented as follows:

<iport> <ivpi> <oport> <ovpi>

The VCI value remains unchanged when cells are switched via a through path. For example, the through path A4 | 10 -> B4 | 20 will map cells received on port A4 with VPI: 10 and any VCI to port B4 with VPI: 20 and the same VCI.

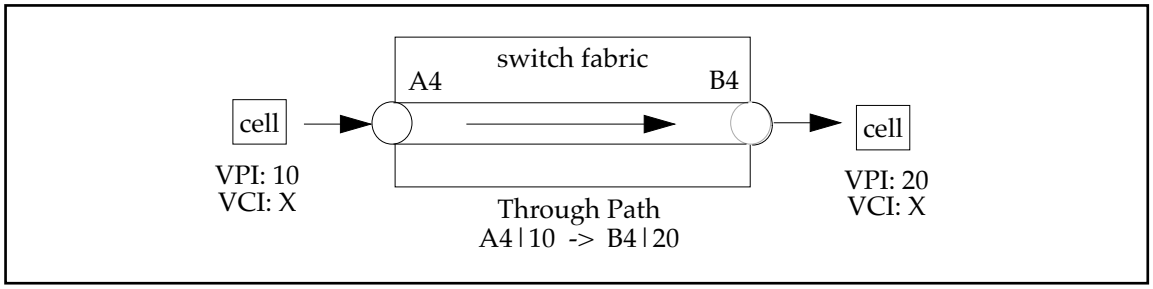

Figure 1.5 - An Example of a Through Path

By definition, through paths will only switch cells in one direction; they are unidirectional. For example, switch fabric X is configured with the through path B1 | 20 -> C1 | 20. If a cell is received on port C1 with VPI: 20, it will not be transmitted on port B1 with a new VPI: 20. In order for this to happen, the through path C1 | 20 -> B1 | 20 must exist as well. Since through paths are unidirectional, two through paths are necessary for bidirectional communication.

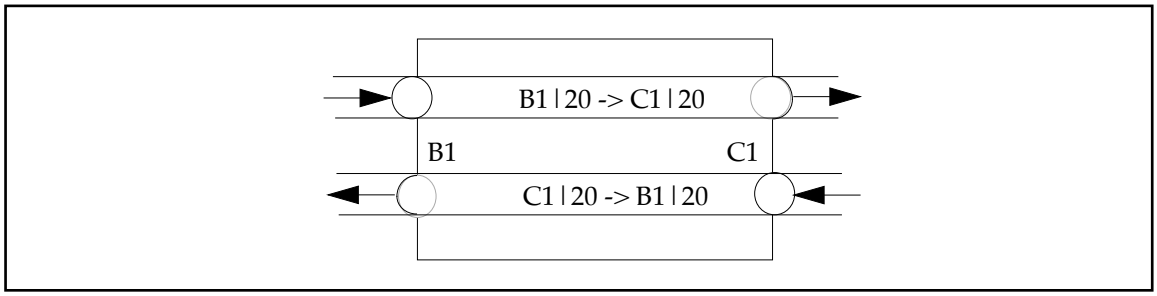

Figure 1.6 - Through Paths are Unidirectional

1.2.2 Originating and Terminating Paths

As previously noted, originating and terminating paths are points at which a virtual path originates and terminates. For example, if a virtual path exists from switch fabric A to switch fabric B, then there must be an originating path on switch fabric A and a terminating path on switch fabric B.

An originating path is defined by two parameters: output VPI and output port. Similarly, a terminating path is defined by the parameters: input VPI and input port. Because originating and terminating paths do not define the way cells are switched through an ATM switch fabric, virtual channels must exist to switch cells from a terminating path to an originating path. (See the section about virtual channels for more information). Originating and terminating paths are the endpoints of virtual paths and are used primarily for bandwidth allocation.

The bandwidth allocated to originating and terminating paths is used to control the amount of bandwidth entering or leaving a virtual path. The total bandwidth used by virtual channels on an originating path or a terminating path cannot exceed the amount of bandwidth allocated to that path. For example, as illustrated in Figure 1.7, if each of the four virtual channels shown is using 10 Mbps of bandwidth, then the originating and terminating paths must have at least 40 Mbps of bandwidth allocated.

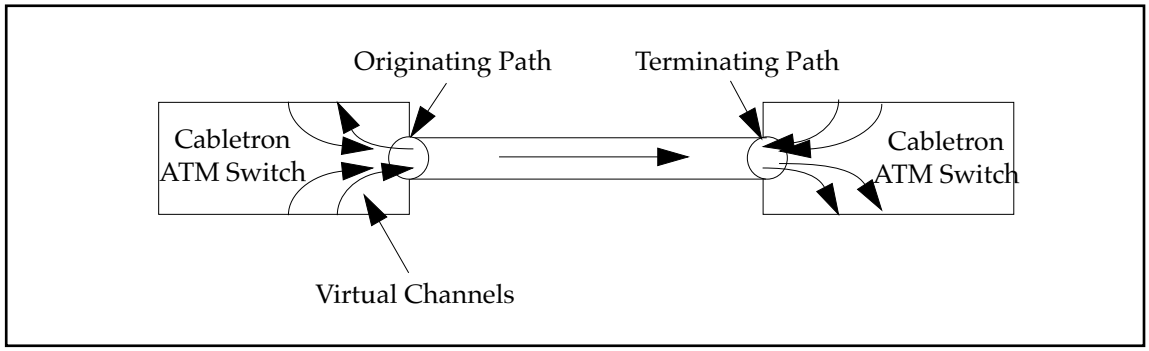

Figure 1.7 - Using Originating and Terminating Paths for Bandwidth Allocation

1.3 Listing Virtual Paths

By logging in to AMI (see Appendix A of this manual for information about logging into AMI), it is possible to display either all of the existing virtual paths on an individual switch fabric or all of the existing virtual paths on a specified port. To list all of the existing virtual paths on an individual switch fabric, enter the following parameters:

configuration vpc show

```
Input       Output
Port  VPI   Port  VPI    MaxBW      BW MaxVCs  VCs  UPC  Prot
1C1     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1C2     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1C3     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1C4     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1D1     0   terminate   155.0M   0.0K    255    2  N/A  pvc
1D2     0   terminate   155.0M   0.0K    255    5  N/A  pvc
1D3     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1D4     0   terminate   155.0M   0.0K    255    4  N/A  pvc
1CTL    0   terminate      N/A   0.0K    511   28  N/A  pvc
originate  1C1     0   155.0M   0.0K    255    4  N/A  pvc
originate  1C2     0   155.0M   0.0K    255    4  N/A  pvc
originate  1C3     0   155.0M   0.0K    255    4  N/A  pvc
originate  1C4     0   155.0M   0.0K    255    4  N/A  pvc
originate  1D1     0     0.0K   0.0K    256    3  N/A  pvc
originate  1D2     0   155.0M   0.0K    255    4  N/A  pvc
originate  1D3     0   155.0M   0.0K    255    4  N/A  pvc
originate  1D4     0   155.0M   0.0K    255    4  N/A  pvc
originate  1CTL    0      N/A   0.0K    511   36  N/A  pvc

Press return for more, q to quit: q
```

The fields in this display are defined as follows:

Input Port — Lists the number of the input port of the virtual path. Lists originate if it is an originating path.

Input VPI — Indicates the input virtual path.

Output Port — Lists the number of the output port of the virtual path. Lists terminate if it is a terminating path.

Output VPI — Shows the output virtual path.

MaxBW Shows the maximum amount of bandwidth (in megabits/second) that is available for the virtual channels using this path.

BW Shows the amount of bandwidth (in megabits/second) that has been reserved for the virtual channels using this path.

MaxVCs Lists the maximum number of virtual channels that may use this originating or terminating path.

VCs Lists the number of virtual channels that are currently using this originating or terminating path.

UPC Shows the integer index that refers to a specific traffic contract assigned to this path.

Protocol Indicates `pvc` for a Permanent Virtual Circuit, `spans` for a Switched Virtual Circuit, or `uni` for a UNI 3.0 signalling path.

To list advanced information about all of the existing virtual paths, enter the following parameters:

configuration vpc show advanced

```
Input       Output
Port  VPI   Port  VPI   Shape VBROB BuffOB
1C1     0   terminate     N/A   N/A    N/A
1C2     0   terminate     N/A   N/A    N/A
1C3     0   terminate     N/A   N/A    N/A
1C4     0   terminate     N/A   N/A    N/A
1D1     0   terminate     N/A   N/A    N/A
1D2     0   terminate     N/A   N/A    N/A
1D3     0   terminate     N/A   N/A    N/A
1D4     0   terminate     N/A   N/A    N/A
1CTL    0   terminate     N/A   N/A    N/A
originate  1C1     0           100    100
originate  1C2     0           100    100
originate  1C3     0           100    100
originate  1C4     0           100    100
originate  1D1     0           100    100
originate  1D2     0           100    100
originate  1D3     0           100    100
originate  1D4     0           100    100
originate  1CTL    0           100    100
```

The fields in this display are defined as follows:

Input Port Lists the number of the input port of the virtual path. Lists originate if it is an originating path.

Input VPI Indicates the input virtual path.

Output Port Lists the number of the output port of the virtual path. Lists terminate if it is a terminating path.

Output VPI Shows the output virtual path.

Shape Indicates whether or not traffic shaping has been enabled for this originating path.

VBROB Lists the bandwidth overbooking level assigned to this path, specified as a percentage. Valid values are integers greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

BuffOB Shows the buffer overbooking level assigned to this path, specified as a percentage. Valid values are integers greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

1.4 Virtual Channels

Virtual channels "ride" inside of virtual paths. The combination of the two specifies a virtual connection. On a switch fabric, each virtual channel switches cells with a specific VPI and VCI received on a specific port to another port with a new VPI and a new VCI. Unlike through paths, virtual channels describe a single virtual connection between two endpoints connected to a switch fabric.

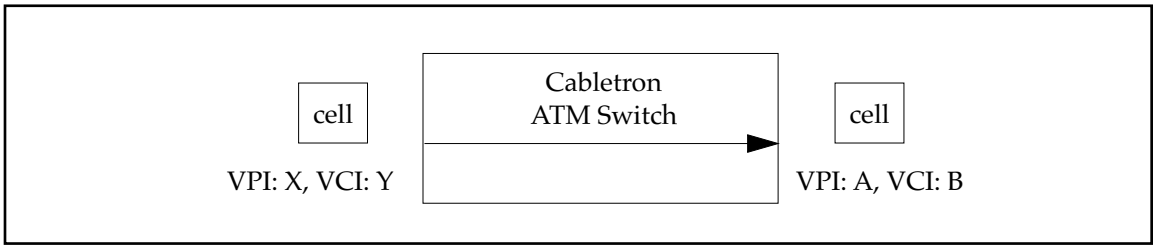

Figure 1.8 - An Example of a Virtual Channel

Six parameters are needed to define a virtual channel: input port, input VPI, input VCI, output port, output VPI, and output VCI. Virtual channels are represented by the following notation:

<iport> <ivpi> <ivci> <oport> <ovpi> <ovci>

Virtual channels switch cells using both the VPI and VCI values. Both the VPI and VCI values may change when a cell is switched via a virtual channel. For example, the virtual channel C2|1|20 -> D2|9|25 will switch cells received on port C2 with VPI: 1 and VCI: 20 such that they are transmitted out port D2 with VPI: 9 and VCI: 25.

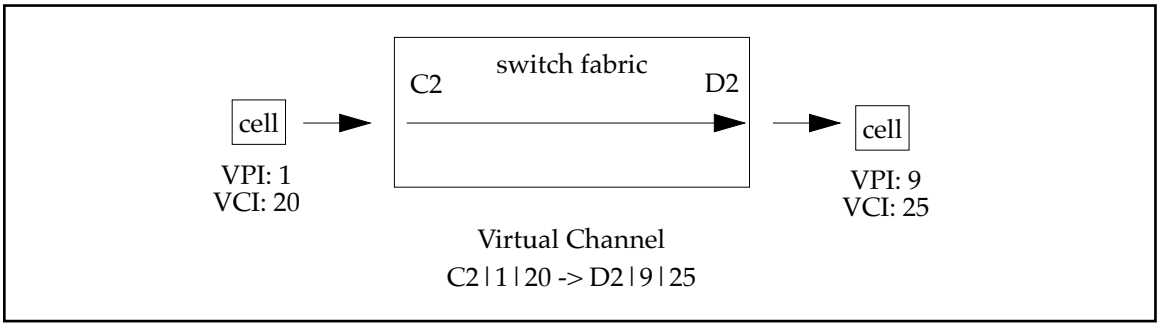

Figure 1.9 - Example of a Virtual Channel

In order to establish two-way communications between two ports on a switch fabric, two virtual channels are necessary because virtual channels are unidirectional. For example, switch fabric A is configured with the virtual channel C3|7|12 -> D1|8|2. If a cell is received on port D1 with VPI: 8 and VCI: 2, it will not be transmitted out port C3 with VPI: 7 and VCI: 12. An additional channel, namely D1|8|2 -> C3|7|12, would have to exist.

Figure 1.10 - Virtual Channels are Unidirectional

Before a virtual channel can be created, the corresponding terminating and originating paths must exist. For example, before the channels shown on the switch fabric in Figure 1.11 can be created, the terminating path C3|3 must exist.

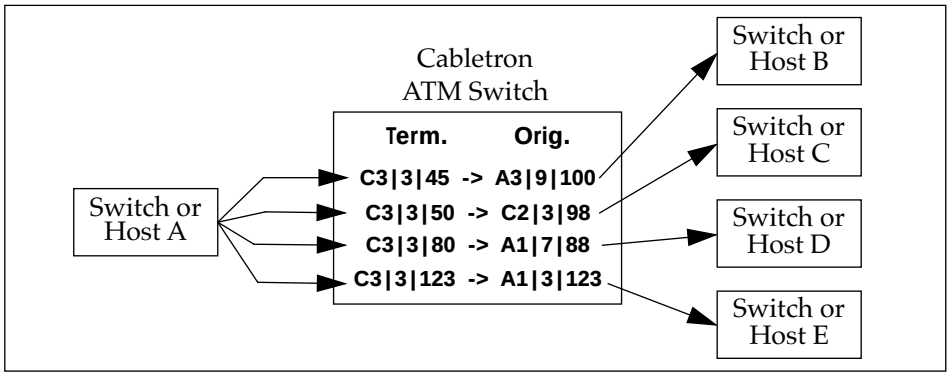

Figure 1.11 - Virtual Channels Created on Terminating Path C3|3

Similarly, before the virtual channels shown in Figure 1.12 can be created, the originating path C2|2 must exist.

.

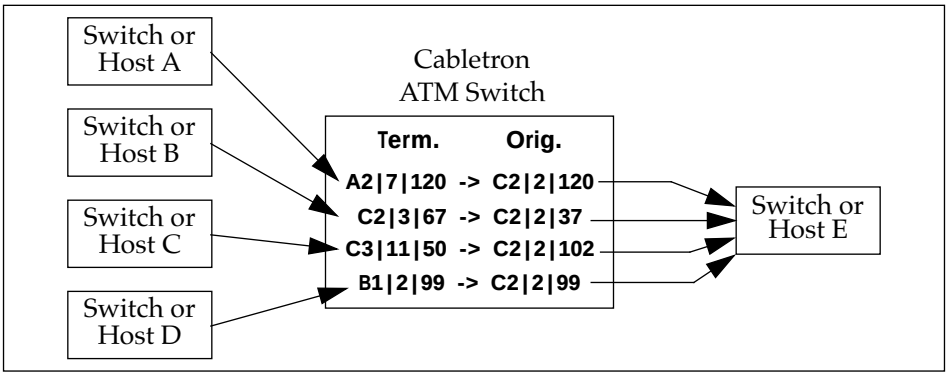

Figure 1.12 - Virtual Channels Created on Originating Path C2|2

Furthermore, in these examples, the terminating path C3|3 and originating path C2|2 must have enough bandwidth allocated to support the total bandwidth used by the virtual channels (Reference Figure 1.7, Using Originating and Terminating Paths for Bandwidth Allocation).

1.4.1 Listing Virtual Channels

By logging in to AMI (see Appendix A of this manual for information about logging into AMI), it is possible to display either all of the existing virtual channels on an individual switch fabric or all of the existing virtual channels on a specified port. To list all of the existing virtual channels on an individual switch fabric, enter the following parameters:

configuration vcc show

```
Input        Output
Port  VPI  VCI  Port  VPI  VCI  UPC  Protocol
1C1     0    5  1CTL    0   34    0  uni30
1C1     0   14  1CTL    0   33    0  spans
1C1     0   15  1CTL    0   32    0  spans
1C1     0   16  1CTL    0   58    0  uni30
1C2     0    5  1CTL    0   37    0  uni30
1C2     0   14  1CTL    0   36    0  spans
1C2     0   15  1CTL    0   35    0  spans
1C2     0   16  1CTL    0   59    0  uni30
1C3     0    5  1CTL    0   40    0  uni30
1C3     0   14  1CTL    0   39    0  spans
1C3     0   15  1CTL    0   38    0  spans
1C3     0   16  1CTL    0   60    0  uni30
1C4     0    5  1CTL    0   43    0  uni30
1C4     0   14  1CTL    0   42    0  spans
1C4     0   15  1CTL    0   41    0  spans
1C4     0   16  1CTL    0   61    0  uni30
1D1     0   14  1CTL    0   45    0  spans

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

Input Port	Displays the incoming port number of the virtual channel.
Input VPI	Shows the incoming virtual path number.
Input VCI	Indicates the incoming virtual channel number.
Output Port	Displays the outgoing port number of the virtual channel.
Output VPI	Shows the outgoing virtual path number.

Output VCI	Indicates the outgoing virtual channel number.
UPC	Shows the integer index that refers to the specific traffic contract assigned to this VCI.
Protocol	Displays what type of protocol is running on this channel. Can be `spans`, `pvc`, or `uni30`.

1.5 Creating Permanent Virtual Circuits (PVCs)

This section contains examples of the steps necessary to create a virtual path or a virtual channel through the network.

NOTE: When these virtual paths and virtual channels are created, a command is entered automatically into the current configuration database, which means that this PVC will be created every time the switch control processor (SCP) is restarted. Note that this configuration database can not be edited. The user can, however, enter AMI and manually delete the PVC that was created. It is recommended that users back up the configuration database frequently.

1.5.1 Creating a Virtual Path

To create a new virtual path, log in to AMI (see Appendix A of this manual for information about logging into AMI), and enter the following parameters:

configuration vpc new <iport> <ivpi> <oport> <ovpi> [-upc <index>]

or

configuration vpc new <port> <vpi> term [-reserved <Kbs>] [-maxvci <maxvci>]

or

configuration vpc new <port> <vpi> orig [-reserved <Kbs>] [-maxvci <maxvci>]

advanced options:

[-shapeovpi <vpi>] [-vbrob <percent>] [-vbrbuffob <percent>]

or

[-shapeivpi <vpi>]

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

-upc <index> Indicates the integer index that refers to a specific traffic contract. If no index is specified, then no traffic policing will take place on this VPI. It is assigned a UPC index of 0, and all traffic on this VPI is treated as UBR traffic. This is the default.

or

port Indicates the incoming port number.

vpi Indicates the incoming virtual path number.

term Specifies the virtual path to be created as a terminating path.

-reserved <Kbs> Indicates the amount of bandwidth specified in Kbps that the user wants to reserve on this path.

-maxvci <maxvci> Indicates the maximum number of channels that can be created on this path.

or

port Indicates the incoming port number.

vpi Indicates the incoming virtual path number.

orig Specifies the virtual path to be created as an originating path.

-reserved <Kbs> Indicates the amount of bandwidth specified in Kbps that the user wants to reserve on this path.

-maxvci <maxvci> Indicates the maximum number of channels that can be created on this path.

The advanced options are as follows:

-shapeovpi <vpi> Indicates the output port of a traffic shaping originating path. Setting this value configures traffic shaping on the originating path. Cells bound for the network leave the traffic shaping port with this value. When the traffic shaping port is the WAN port, this value equals the input VPI of the through path form the traffic shaping port to the WAN port.

-vbrob <percent> Indicates the bandwidth overbooking level assigned to this path, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

-vbrbuffob <percent> Indicates the buffer overbooking level assigned to this path, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

-shapeivpi <vpi> Indicates the incoming VPI. This can only be set when creating a through path. When the traffic shaping port is not the port connected to the WAN, a through path must be created from the WAN port to the traffic shaping port. Cells arrive from the network at the traffic shaping port with this value identical to the VPI of the terminating path at the traffic shaping port.

NOTE: Traffic can be shaped only on originating paths on Series C network module ports.

The following is an example of how to create a virtual path on an SFCS-1000. To create a through path going in port 2A1, vpi 1 on on the switch board installed in slot 2 and going out port 4B1, vpi 1 on on the switch board installed in slot 4, enter the following:

```
localhost::configuration vpc> new 2a1 1 2e4 1
localhost::configuration vpc> new 2e4 1 2a1 1

localhost::configuration vpc> new 4b1 1 4e2 1
localhost::configuration vpc> new 4e2 1 4b1 1
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

At the same time, a command is entered automatically into the current configuration database, which means that this virtual path will be created every time the SCP is restarted.

1.5.1.1 Terminating a PVC at a Switch

Sometimes it is necessary to create a PVC between a host and a switch fabric that is at a remote location. In this case, the PVC should be created from the host to the control port (CTL) of the SCP and vice versa.

Some additional configuration is necessary for communication to be established between the host and the switch fabric. The SCP needs an entry in its ATM ARP cache in order to send cells destined for the host with the correct VPI and VCI and to pass received cells with a specific VPI and VCI to IP. This configuration can be done using AMI as shown in the following subsections. (See the `atmarp (8c)` man page for more information.)

1.5.1.2 Creating ATM ARP Entries

To create a FORE IP PVC ARP entry, log in to AMI (see Appendix A of this manual for information about logging into AMI). Data on this PVC is encapsulated using null encapsulation (also known as VC-based multiplexing) as specified in RFC-1483. Enter the following parameters:

configuration atmarp newforeip <host> <vpi> <vci> (4|5) [<interface>]

host	Indicates the IP address of the remote host.
vpi	Indicates the virtual path number. This must be 0.
vci	Indicates the virtual channel number.
4|5	Designates the connection's ATM Adaptation Layer (AAL) type.
interface	Indicates the FORE IP interface to be used. The default is `asx0`.

Once the parameters are entered, the FORE IP PVC ARP entry is created instantly by the SCP. At the same time, a command is entered automatically into the current configuration database which creates this ATM ARP entry every time the SCP is restarted.

To create a new Classical IP PVC ARP entry, log in to AMI. All data is sent LLC/SNAP encapsulated. Enter the following parameters:

configuration atmarp newclassicalip <host> <vpi> <vci> [<interface>]

host	Indicates the host IP address of the remote IP endstation.
vpi	Indicates the virtual path number of the Classical IP PVC.
vci	Indicates the virtual channel number of the Classical IP PVC.
interface	Indicates the name of the Classical IP interface to be used for this connection. The default is `qaa0`.

Once the parameters are entered, the Classical IP PVC ARP entry is created instantly by the SCP. At the same time, a command is entered automatically into the current configuration database which creates this ATM ARP entry every time that the SCP is restarted.

1.5.1.3 Listing ATM ARP Entries

To verify that the ARP entries exist correctly for the outgoing PVC connection from the SCP to the host, log in to AMI (see Appendix A of this manual for information about logging into AMI). To display the ATM ARP cache, enter the following parameters:

configuration atmarp show

```
IPaddress        If    VPI   VCI   AAL    Type              Direction
198.29.22.9      asx0  0     63    aal5   foreIpSVC         pending
198.29.22.15     asx0  0     231   aal5   foreIpSVC         pending
198.29.22.37     asx0  0     65    aal34  foreIpSVC         pending
IPaddress        If    NSAP Address
198.29.17.3      qaa0  0x47.0005.80.ffe100.0000.f21b.0138.002048102754.00
198.29.17.10     qaa0  0x47.0005.80.ffe100.0000.f21b.0137.002048100be6.00
198.29.17.15     qaa0  0x47.0005.80.ffe100.0000.f21b.0137.00204810048d.00
198.29.17.52     qaa0  0x47.0005.80.ffe100.0000.f21b.0138.0020481b0138.00
```

The fields in this display have the following meanings:

IPaddress	Indicates the IP address for this connection.
If	Shows the name of the IP interface for this connection.
VPI	Displays the virtual path number.
VCI	Displays the virtual channel number.
AAL	Displays the AAL type of the given connection.
Type	Lists the kind of connection. Can be foreIpPVC, foreIpSVC, classicIpPVC, or classicIpSVC.
Direction	Pending means that a connection has not (yet) been established. Incomplete means that the IP-to-ATM address mapping is not yet known for the given IP address.
NSAP Address	Lists the NSAP address for this connection.

1.5.2 Creating a Virtual Channel

To create a new virtual channel, log in to AMI and enter the following parameters:

configuration vcc new <iport> <ivpi> <ivci> <oport> <ovpi> <ovci> -upc <index>

iport	Indicates the incoming port number.
ivpi	Indicates the incoming virtual path number.
ivci	Indicates the incoming virtual channel number.
oport	Indicates the outgoing port number.
ovpi	Indicates the outgoing virtual path number.
ovci	Indicates the outgoing virtual channel number.
-upc<index>	Indicates the integer index that refers to a specific traffic contract. If no index is specified, then no traffic policing will take place on this VCI. It will be assigned a UPC index of 0, and all traffic on this VCI will be treated as UBR traffic. This is the default.

The following is an example of how to create a virtual channel on an SFCS-1000. To create a vcc going in port 2A1, vpi 0, vci 100 on the switch board installed in slot 2 and going out port 4B1, vpi 0, vci 100 on the switch board installed in slot 4, enter the following:

```
localhost::configuration vcc> new 2a1 0 100 2e4 0 100
localhost::configuration vcc> new 2e4 0 100 2a1 0 100

localhost::configuration vcc> new 4b1 0 100 4e2 0 100
localhost::configuration vcc> new 4e2 0 100 4b1 0 100
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

Once the parameters are entered, the virtual channel will be created instantly by the SCP. At the same time, a command is entered automatically into the current configuration database, which means that this virtual channel will be created every time the SCP is restarted.

1.6 Traffic Types

Quality of Service (QOS) Management is based on the bandwidth parameters associated with a virtual connection and the class of service and ATM Adaptation Layer (AAL) used for that connection. In order to support voice, video, and data, the ATM Forum has defined four classes of service, or traffic types: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR), and Unspecified Bit Rate (UBR).

- At connection set-up time, traffic that uses a CBR parameter, such as a voice signal, makes a request for a dedicated Peak Cell Rate (PCR). Once the amount is defined, the ATM network must be able to guarantee that amount for the duration of the connection.
- At connection set-up time, traffic that uses a VBR parameter, such as a video and data, makes a request for a dedicated PCR, Sustainable Cell Rate (SCR), and Maximum Burst Size (MBS). Once these amounts are defined, the ATM network must be able to guarantee these rates for the duration of the connection.
- Connections that use an ABR parameter, such as pure data, must adjust to changing network conditions and can only use the bandwidth that is available at the time after the CBR and VBR traffic is serviced. This traffic is the most difficult to accommodate since the switches that carry it must constantly update the endpoints as to the amount of available bandwidth at a given point in time.
- UBR traffic, such as broadcast information and ARP messages, is also known as “best effort” service. UBR provides no bandwidth guarantees.

Because ATM is designed to provide a single network to transport this variety of traffic classes, Cabletron’s traffic policing and Connection Admission Control (CAC) schemes are vital to allowing this mix of traffic to flow smoothly.

1.7 Traffic Policing (Usage Parameter Control)

Traffic policing, also known as Usage Parameter Control (UPC), is a method of assessing the cells entering the switch for conformance with pre-established traffic bandwidth contracts. Those cells that exceed the specified contract are "tagged" or "dropped," depending on what is defined in the contract. This ensures that the connections with reserved bandwidth are not exceeding their reservations. Cabletron Systems' switches use a combination of "leaky bucket," or Generic Cell Rate Algorithm (GCRA) hardware in the switch fabric and user-configurable parameters in AMI to perform these policing functions.

1.7.1 Setting the CLP Bit

First, it is important to understand the concept of tagging and dropping. Each ATM cell has a Cell Loss Priority (CLP) bit which indicates if the network can drop it under congested conditions. When the CLP bit is set to 0 (or CLP = 0), the cell will be assessed for compliance with traffic parameters. If the traffic parameters dictate that non-compliant cells should be "tagged," the CLP bit will be set to 1 (or CLP = 1), which means that upon experiencing congestion further in the network, these CLP = 1 cells will be dropped (because you cannot tag them again).

1.7.2 Leaky Bucket Algorithm

The next important concept is the leaky bucket algorithm. Leaky buckets are a mechanism by which cells entering the switch fabric are monitored for compliance with UPC traffic contracts that have been negotiated at connection set-up time. Before the leaky buckets are discussed, it is important to understand the parameters that are being measured by the buckets. They are as follows:

- Peak Cell Rate (PCR) - the maximum number of cells per second.
- Cell Delay Variation Tolerance (CDVT) - the tolerance for variation in the inter-arrival time of these cells, or the amount of jitter that can be accepted by the network.
- Sustainable Cell Rate (SCR) - the average rate of cell transmission for this connection, taking bursting into account.
- Maximum Burst Size (MBS) - the maximum amount of cells that can be transmitted at the PCR.

The leaky bucket algorithm is a timer which measures the cells entering the switch fabric against the parameters listed above. As a cell arrives, the timer assesses if the cell is on time, late, or early. If the cell is determined to be on time or late (within the traffic parameters specified), the cell is allowed to pass and no changes are made to its CLP bit. If the cell is early (which exceeds the specified parameters and would create congestion), the cell is either dropped or tagged (the CLP bit is set to 1), depending on the specified contract.

The first bucket in this analogy measures the PCR, or the rate at which the bucket drains. It also considers the CDVT, or the depth of the bucket. The second bucket measures the SCR, or the rate at which the bucket drains, and the MBS, or the depth of the second bucket.

1.7.3 UNI 3.0 UPC Traffic Contract Parameters

The ATM Forum has defined different types of traffic contracts to be used in conjunction with these leaky buckets. The parameters that make up these types of contracts are defined as follows:

- pcr0 - measures PCR for cells with CLP = 0
- pcr01 - measures PCR for cells with CLP = 0 + cells with CLP = 1 (all cells)
- scr0 - measures SCR for cells with CLP = 0
- scr01 - measures SCR for cells with CLP = 0 + cells with CLP = 1 (all cells)
- mbs0 - measures MBS for cells with CLP = 0
- mbs01 - measures MBS for cells with CLP = 0 + cells with CLP = 1 (all cells)
- tag - sets CLP bit = 1 for non-compliant CLP = 0 cells

The specific combinations of these parameters that make up the ATM Forum contracts are defined as follows:

1. ubr
2. cbr <pcr01>
3. cbr0 <pcr0> <pcr01> [tag]
4. vbr <pcr01> <scr01> <mbs01>
5. vbr0 <pcr01> <scr0> <mbs0> [tag]

The `ubr` contract is for UBR traffic. Since this is best-effort traffic with no bandwidth guarantees provided, this type of traffic can not be policed against bandwidth parameters.

The `cbr <pcr01>` contract is for CBR traffic. It only uses the first leaky bucket to assess the PCR of the combination of CLP = 0 cells plus the CLP = 1 cells. All non-compliant cells are dropped.

The `cbr0 <pcr0> <pcr01> [tag]` contract is for CBR traffic. It uses the first leaky bucket to assess the PCR of the combination of CLP = 0 cells plus the CLP = 1 cells and to assess the group of CLP = 0 cells separately. All non-compliant cells in both groups are dropped if the tag option is not set. If the tag option is set, the non-compliant CLP = 0 cells in both groups are tagged and the non-compliant CLP = 1 cells are dropped.

The `vbr <pcr01> <scr01> <mbs01>` contract is for VBR traffic. It uses both leaky buckets simultaneously. The first bucket assesses the PCR of the combination of CLP = 0 cells plus the CLP = 1 cells and the second bucket assesses the SCR and MBS of this same combination. All non-compliant cells are dropped even if they comply with one bucket.

The `vbr0 <pcr01> <scr0> <mbs0> [tag]` contract is for VBR traffic. It uses both leaky buckets simultaneously, as well. The first bucket assesses the PCR of the combination of CLP = 0 cells plus the CLP = 1 cells, but the second bucket assesses the SCR and MBS of CLP = 0 cells only. All non-compliant cells are dropped even if they comply with one bucket, provided that the tag option is not set. If the tag option is set, the non-compliant CLP = 0 cells in all groups are tagged and the non-compliant CLP = 1 cells are dropped.

1.7.4 AMI UPC Commands

AMI allows the user to create a UPC contract using these combinations of traffic parameters. To create a UPC contract, log into AMI and enter the following parameters:

```
localhost::configuration upc> new <index> [<UPC>] [-cdvt <us>] [aal5epd] [-name <name>]
Where UPC is one of the following combinations of traffic parameters:
        ubr
        cbr  <pcr01>
        cbr0 <pcr0> <pcr01> [tag]
        vbr  <pcr01> <scr01> <mbs01>
        vbr0 <pcr01> <scr0> <mbs0> [tag]
```

The <index> parameter is an arbitrary, user-assigned integer that represents this particular contract. This number will be entered as a parameter when creating a VPC, VCC, or SPANS signalling path that you want this contract to assess.

> **NOTE:** Remember, when you create this UPC contract, it is not actually used until you assign it to a VPC, VCC, or SPANS path.

The [<UPC>] parameter is one of the five combinations of traffic parameters outlined in Section 1.7.3.

The [-cdvt <us>] parameter is an optional CDVT rate assigned by the user and specified in microseconds.

The [aal5epd] parameter is also optional. If entered, Early Packet Discard (EPD) is enabled on this connection. If it is not entered, then EPD is not enabled on this connection.

The [-name <name>] parameter is an optional, user-assigned name that helps the user remember for what kind of connection this particular contract should be used.

For example, to create a contract for VBR traffic in which the tagging option is used, you would enter the following:

```
localhost::configuration upc> new 1 vbr 1000 500 100 tag -name video-vbr
```

CHAPTER 2 Configuring a Classical IP ATM Network

2.1 Introduction

This chapter describes how to design, configure, and maintain a Classical IP ATM network. The term classical indicates that the ATM network has the same properties as existing legacy LANs. That is, even though ATM technology allows for large, globally connected networks, for example, it is only used in the LAN environment as a direct replacement of existing LAN technology. The classical model of LANs connected through IP routers is maintained in ATM networks. RFC-1577 provides the standard for Classical IP over ATM.

Classical IP over ATM is different than IP in legacy LANs in that ATM provides a virtual connection environment through the use of Permanent Virtual Circuits (PVCs) and/or Switched Virtual Circuits (SVCs). SVC management is performed via the ATM Forum UNI 3.0 Specification, which specifies Q.2931. Q.2931 is a broadband signalling protocol designed to establish connections dynamically at the User-Network Interface (UNI). Q.2931 uses Service Specific Connection Oriented Protocol (SSCOP) as a reliable transport protocol, and all signalling occurs over VPI: 0, VCI: 5. Q.2931 connections are bidirectional, with the same VPI/VCI pair used to transmit and receive.

Once a Classical IP connection has been established, IP datagrams are encapsulated using IEEE 802.2 LLC/SNAP and are segmented into ATM cells using ATM Adaptation Layer type 5 (AAL5). In addition, the default Maximum Transmission Unit (MTU) is 9,180 bytes (the SNAP header adds 8 more bytes) with a maximum packet size of 65,535 bytes.There is currently no support for IP broadcast datagrams or IP multicast datagrams in a Classical IP environment.

2.1.1 Logical IP Subnets

An important concept in Classical IP networks is that of a Logical IP Subnet (LIS). An LIS is a group of hosts configured as members of the same IP subnet (that is, they have the same IP network and subnetwork numbers). In this sense, one LIS can be equated to one legacy LAN. It is possible to maintain several overlaid LISs on the same physical ATM network. Therefore, in a Classical IP ATM network, placing a host on a specific subnet is a logical choice rather than a physical one. In this type of environment, communication between hosts in different LISs is only permitted by communicating through an IP router which is a member of both LISs (as per RFC-1577).

The number of LISs, and the division of hosts into each LIS, is purely an administrative issue. Limitations of IP addressing, IP packet filtering, and administrative boundaries may guide a manager into establishing several LISs onto a single ATM network. Keep in mind, though, that communication between LISs must occur through IP routers.

2.1.2 Classical IP Interfaces

In order to support routing between multiple LISs, the host adapter software allows a host to be configured as a member of (and a router between) up to four distinct LISs per physical ATM interface. Each LIS membership is through a separate Classical IP network interface. Existing system level IP routing configuration tools are used to control routing through each of the Classical IP interfaces in the same manner as routing among several physical interfaces. Note that even though each Classical IP interface associated with a given physical interface uses the same physical hardware, they are each configured separately with their own MTU, IP address, and ATM address.

The name of each of the Classical IP interfaces begins with `qa`. All of the Classical IP interfaces associated with physical unit zero will have `a` as the next letter. All of the Classical IP interfaces associated with physical unit one will have `b` as the next letter, and so forth. Finally, each Classical IP interface has its interface number as a suffix. As an example of the above naming convention, the name of the third Classical IP interface (unit 2) on physical unit one is `qab2`.

2.1.3 SPANS Interface

While each of the `qa` interfaces for a given physical interface is designed to support Classical IP using Q.2931 signalling, a SPANS interface (usually called `fa`, but user configurable) also exists for each physical interface. The `fa` interface supports FORE IP on top of SPANS signalling. FORE IP allows communication using AAL4 or AAL5 with no encapsulation, uses a broadcast ARP for SPANS address resolution, and supports direct communication of all hosts on a physical ATM network without the use of IP routers. Since SPANS and Q.2931 signalling use different VCIs, a host can simultaneously support FORE IP over SPANS as well as Classical IP over Q.2931 on the same physical interface.

As a result of standard IP routing, all traffic sent out an `fa` interface will use FORE IP, while all traffic sent out a `qa` interface will use Classical IP. Each of the `fa` interfaces should be assigned an IP address on a subnet different than the subnets of any of the `qa` interfaces. It is permissible to place multiple `fa` interfaces on the same subnet, and the driver will load balance connections across these interfaces.

It is only necessary to configure the `fa` and `qa` interfaces if the specific service provided by that interface is required. A host sending only Classical IP would not need to configure the `fa` interfaces. Likewise, a host sending only FORE IP would not need to configure the `qa` interfaces. Both the `fa` and `qa` interfaces may be configured simultaneously, but they must be in separate subnets. Remember that Classical IP specific configuration changes can only be done with the `qa` devices, while SPANS specific configuration changes can only be done with the `fa` devices.

2.2 Address Registration and ILMI

Before a host can establish connections over a physical interface, the host must know the NSAP address for that interface. The primary purpose of Interim Local Management Interface (ILMI) is to discover and register these NSAP addresses dynamically.

2.2.1 NSAP Addresses

For private ATM networks, addresses uniquely identify ATM endpoints. The UNI 3.0 address format is modeled after that of an OSI Network Service Access Point, hence the name NSAP address.

Three address formats have been specified: DCC, ICD, and E.164. Cabletron implements the ICD ATM format. Per the UNI 3.0 specification, all private networks should accept initial call set-up messages containing ATM addresses with any of the approved formats and forward the calls as necessary.

An NSAP address consists of the following:

- a 13-byte network-side prefix - The prefix is the NSAP prefix of the SCP to which the host is attached.
- a seven-byte user-side part - This consists of the following:
  - a six-byte End System Identifier (ESI) - The ESI is the unique IEEE MAC address of the interface.
  - a one-byte selector - Although each Classical IP interface for a given physical interface uses the same prefix and ESI, the selector field is the part that indicates the number of the specific Classical IP interface. The selector field is 00 for qaa0, 01 for qaa1, 02 for qaa2, and 03 for qaa3.

2.2.2 Operating with ILMI Support

Cabletron Systems switches running software version 3.0 or greater provide support for ILMI. If ILMI is supported on all of the switches and hosts in a given network, when a switch control processor (SCP) boots up, ILMI enables the SCP to discover all of the hosts attached to it and to send its NSAP prefix associated with the port to those hosts dynamically. In return, the host prepends that prefix to its ESI and selector fields, forming a complete NSAP address. The host then notifies the SCP of its complete NSAP address. These registration SNMP messages are sent and received over AAL5 using VPI: 0, VCI: 16. Once ILMI registration has been completed, then connection setup can occur.

If a host changes network ports after an NSAP address has been registered for its interface, all existing connections will be closed. If the new port is on a different SCP, a new NSAP address (with a different network address prefix) will be registered. The host can then begin to establish new connections.

2.2.3 Operating without ILMI Support

If ILMI is not supported on a particular SCP or host in a given network, then the NSAP addresses must be manually configured. Because a non-Cabletron switch does not support ILMI, it can not supply an NSAP prefix to the hosts. Therefore, the user must assign a unique, valid prefix to the switch. Additionally, the same prefix should be used for all hosts in the LIS.

On the host, `atmarp -n (8c)` is used to configure the NSAP address for a specific interface. The SCP directly attached to this interface is then informed of this NSAP address/port combination through commands in the ATM Management Interface (AMI). Once the host and network have both been informed of this NSAP address/port pair, the host may begin signalling.

Once an interface's NSAP address has been set, it is possible to change it by first running `atmarp -n` to set the new address, and then bringing the interface down and back up (`ifconfig qaXX down`; `ifconfig qaXX up`) for the address change to take effect.

2.2.4 Configuration

The choice to use ILMI for address registration is made at software installation time. Since ILMI uses SNMP as its management protocol, the use of ILMI is tied into `snmpd`. The choice can be made to run FORE's SNMP agent and use ILMI (`snmpd`), run FORE's SNMP agent without using ILMI (`snmpd -n`), or just use ILMI (`snmpd -i` or `ilmid -i`). The commands to run the chosen daemon, or to configure a host's NSAP address statically when not running ILMI, are placed in the firmware download script.

2.3 ARP and ARP Servers

2.3.1 Theory

In order for a host to establish a connection to another host, it must first determine the other host's NSAP address. ATM ARP (ATM address resolution protocol) is the procedure used to resolve an IP address into an ATM address. Since the ATM standards do not currently support broadcast on an ATM LAN, address resolution is performed by direct communication with a special ARP server, rather than broadcasting ARP requests as is done in legacy LANs. Each LIS must have only one ARP server configured, but a single ARP server can be the server for several LISs.

Each host in an LIS must be configured with the NSAP address of the host providing ARP service for its LIS. On a host ARP server, the NSAP address of the ARP server can be obtained by running `atmarp -z` (remember to use the interface associated with the given LIS). The ARP server address is normally configured into each host at installation time (the appropriate commands are placed in the firmware download script), but it may be done at any time by running `atmarp -p`. If the ARP server address is configured in the firmware download script, it will be persistent across reboots, but if it is configured using `atmarp -p`, it will not be persistent across reboots.

Since only one ARP server can be functioning at a time in a given LIS, and since the ARP server's address is manually configured into each host, it is not possible to use multiple, redundant ARP servers to improve robustness. If an ARP server becomes nonfunctional, a new ARP server must be configured, and then each host within the LIS must be configured to use the new ARP server. To configure a new ARP server address on a host, run `atmarp -p` to set the new address, then bring the appropriate interface down and back up (`ifconfig qaXX down`; `ifconfig qaXX up`) for the change to take effect.

2.3.2 Configuring a Host to be an ARP Server

The following procedures list the required steps to configure a host as an ARP server. Before proceeding, several preconditions MUST be met. You must also obtain both the interface name and ARP server NSAP address.

The preconditions are as follows:

1. A FORE adapter must be installed in the network.
2. The FORE Systems software has been installed and the system has been successfully booted.
3. The network must contain devices that recognize ILMI and RFC-1577, such as FORE switches and adapters.
4. When the host was initially configured, Classical IP should have been specified. By specifying Classical IP, ILMI and RFC-1577 are automatically active in the Rev. 3.0 adapter code.

ARP server/client configurations are run on the Classical IP interfaces. Classical IP interface names are of the form `qaaN`, where N is from 0 to 3. Once the prerequisites have been met, the steps for configuring a host as an ARP server are as follows:

1. If not already done during installation, configure the IP, netmask, and broadcast address for the interface using the `ifconfig` command as follows:

`ifconfig qaa0 <ip address> netmask <mask> broadcast <broadcast_address>`

2. Using the host interface name from the previous step, determine the NSAP address of the interface name you wish to make the ARP server with the following command (assuming the interface name is `qaa0`):

`/usr/etc/fore/atmarp -z qaa0`

A typical response is shown below:

```
NSAP addr for qaa0 is 0x47.0005.80.ffe100.0000.f215.0ce5.002048100c47.03
```

NOTE: If the response is all zeros, this indicates that the NSAP address has not been assigned. Check the other `qaaN` interfaces to see if any have been configured. If not, then manually assign an NSAP address as in step 1 above.

If you are using a windowed graphical user interface, keep the window with the NSAP address open because you will need to cut and paste this address into several locations. If not, carefully write down the address so you can add it correctly to the other locations. If you are not registered with ILMI, you will need to assign or manually configure an NSAP address to the interface.

With the preconditions met and the required information obtained, ARP servers may be created with the following steps. Two methods are shown. An ARP server can be configured for a one-time use, which is deactivated if the ARP server is rebooted. This method is shown first. An ARP server can also be permanently configured to be active on boot. This method is shown second.

2.3.2.1 Configuring a Host as an ARP Server for One-Time Use

1. On the host you wish to configure as an ARP server, log in as `root`.
2. On the host, change to the directory in which the ATM boot file is stored. To determine the location of this directory, use your system specific `find` command and search for the location of the `atmarp` utility. The ATM boot files are stored in this directory. (See the list below for the names of the system specific boot files.)

3. From the command prompt in this directory, issue the following command:

atmarp -p <NSAP address> <interface name>

The location of the ATM boot file is a function of the type of host in which the FORE adapter has been installed. On your system, the files may have been copied to another location during the installation process. The following list shows the default locations:

- Solaris
  `/etc/rcS.d/S99rc.sba200`
- SunOS
  `/usr/etc/fore/rc.sba200`
- IRIX
  `/etc/init.d/fore_atm`
- HP-UX
  `/etc/fore_atm`

4. On each host requiring service from the ARP server, use the following command:

atmarp -p <ARP Server NSAP address> <interface name>

2.3.2.2 Configuring a Host as a Permanent ARP Server

To configure an ARP server so that it persists across reboots, use the following procedure:

1. On the host you wish to configure as an ARP server, log in as `root`.
2. In the directory in which the ATM boot file is stored, edit the ATM boot file. Note where this file is in the directory structure from the list of operating systems above. For example, for an ESA-200 adapter, the file name is `fore_atm`. In this file, add the following line:

atmarp -p <NSAP_address> <interface_name>

where the correct NSAP address and interface names are used.

3. On each host requiring service from the ARP server, in the directory in which the ATM boot file is stored, edit the boot file. For example, for an ESA-200 adapter, the file name is `fore_atm`. In this file, add the following line:

`atmarp -p <ARP_server_NSAP_address> <interface_name>`

where the correct NSAP address and interface names are used. For example:

```
atmarp -p 0x47.0005.80.ffe100.0000.f215.0ce5.002048100c47.03 qaa0
```

Once these changes have been made, the ARP server will be active on boot. Similarly, each host attached to the ARP server can access the ARP server on boot.

2.3.2.3 Configuring a Hewlett-Packard Computer as an ARP Server

On a Hewlett-Packard host, use the program `/usr/etc/fore/config_atm` to configure the host as either being connected to the ARP server or being the ARP server itself. Prior to running `config_atm`, see Section 2.3.2 for instructions on how to determine the host NSAP address and interface names to be used. The discovery process is the same for all computers.

2.3.3 Configuring a Cabletron Switch to be an ARP Server

Cabletron's ATM switches also have the capability of being an ARP server. To configure a Cabletron ATM switch as an ARP server, perform the following steps on only one of the SCPs:

1. On one of the SCPs, determine the NSAP address of that SCP for the relevant interface (qaa0 -> qaa3) using the following AMI command:

configuration atmarp> getnsap <interface>

For example:

configuration atmarp> getnsap qaa0

```
qaa0 NSAP address: 47000580ffe1000000f12400de0020481900de00
```

2. Set the NSAP address of the ARP server to be the NSAP address of the interface that you displayed in step 1 using the following AMI command:

configuration atmarp arpserver> set <NSAPaddress> [<interface>]

For example:

configuration atmarp arpserver> set 0x47000580ffe1000000f12400de0020481900de00 qaa0

2.3.4 Classical IP Operation

Once a host knows its own ATM address and the ATM address of its ARP server it will attempt to establish a connection to the ARP server, which will be used to send ARP requests and receive ARP replies. When the connection to the ARP server has been established, the ARP server sends an inverse ARP (InARP) request on the new VC to learn the host's IP address. When an InARP reply is received, the ARP server places that host's IP address to ATM address mapping in its ARP cache. Therefore, over time, the ARP server dynamically learns the IP-to-ATM address mappings of all the hosts in its LIS. It can then respond to ARP requests directed toward it for hosts in its LIS.

NOTE: In order for a host to communicate with an ARP server, it must have learned its own ATM address and have been configured with the ATM address of the ARP server.

A host can not resolve the ATM addresses of hosts in its LIS unless it can communicate with its ARP server.

Since there is no mechanism for ARP servers to exchange mapping information with each other, it is imperative that each LIS be configured with only one ARP server.

When a host wants to communicate with another host in its LIS, it first sends an ARP request to the ARP server containing the IP address to be resolved. When an ARP reply is received from the ARP server, the host creates an entry in its ARP cache for the given IP address and stores the IP-to-ATM address mapping. This ARP cache entry will be marked as complete. To ensure that all of the IP-to-ATM address mappings known by a certain host are up-to-date, hosts are required to age their ARP entries. A host must validate its ARP entries every 15 minutes (20 minutes on an ARP server).

A host validates its SVCs by sending an ARP request to the ARP server. A host validates its PVCs, and an ARP server validates its SVCs, by sending an InARP request on the VC. If a reply is not received, the ARP entry is marked invalid. Once an ARP entry is marked invalid, an attempt is made to revalidate it before transmitting. Transmission will proceed only when validation succeeds. If a VC associated with an invalid ARP entry is closed, the entry is removed.

2.3.5 Operational Issues

Certain hosts in an LIS may not support Classical IP. It is still possible to communicate with these hosts (and for these hosts to communicate with one another) by using static ARP entries. If a host does not support Classical IP, its IP-to-ATM address mapping should be placed in its ARP server's cache as a static entry. This allows other hosts that do support Classical IP to contact their ARP server as usual and obtain the correct address mapping. If a host that does not support Classical IP wants to initiate connections, the IP-to-ATM address mappings of the destination hosts should be put in its ARP cache, again as static entries. By using static ARP entries in the above fashion, the ability for all hosts to communicate is maintained.

There are some restrictions on the number of hosts that can be maintained as static ARP entries. They are as follows:

- In the default configuration, a host can only have approximately 250 virtual connections open simultaneously. This means that an ARP server can only serve 250 clients, since each client must maintain a connection with its ARP server. This may be a limitation if the ARP server is servicing multiple LISs.
- On the host, it is possible to increase the number of connections allowed by first using `cpath (8c)` to delete path 0 for the given port, and then using `cpath` to recreate path 0 for the given port, but using the `-v` option to `cpath` to specify the number of connections to allow over that path.
- Hosts support a maximum of 1,024 connections.

2.4 Classical IP PVCs

2.4.1 Theory and Configuration

Normally, ATM connections in a Classical IP environment are established dynamically using Q.2931. ARP, ILMI, and Q.2931 all work together as described previously to set up an SVC. If a host from another vendor does not support Classical ARP or ILMI, it is still possible to set up an SVC using work-arounds. If a host or a switch in an LIS does not support Q.2931, however, it is not possible to establish an SVC. In this case, a Classical IP PVC can be used for communication.

On each of the hosts, `atmarp -c` is used to establish the PVC. An unused VPI/VCI pair must be chosen for each host. PVCs using the chosen VPI/VCI pairs must also be set up from each of the hosts to their connecting switch, and then on all of the switches between the two connecting switches.

NOTE: Both the incoming and outgoing connections are set up simultaneously on the host, but they must be set up individually on the switches. The same VPI/VCI pair is used by a host to send on the PVC as well as receive on the PVC. The IP datagrams are sent over the PVC using AAL5 with LLC/SNAP encapsulation.

2.4.2 Revalidation and Removal

Normally, the device driver periodically checks that its PVCs are still established and functioning. A host revalidates a PVC every 15 minutes by sending InARP requests over the PVC, if the user specifies that revalidation should occur by choosing 1 as the revalidate option to `atmarp -c`. If the equipment attached to the Cabletron equipment supports revalidation, the user must choose 1 as the revalidate option to `atmarp -c`. If an InARP reply is not received, the revalidation fails, the PVC is marked invalid (as shown through `atmarp -a`), and communication over the PVC is no longer possible.

Once a PVC is marked invalid, an attempt is made to validate the PVC before transmitting. Transmission will proceed only when validation succeeds. It is possible to disable this revalidation feature by specifying 0 as the revalidate option to `atmarp -c`. This is often desirable when the remote end of the PVC (such as a video camera) does not support InARP.

A Classical IP PVC is removed on the host side using `atmarp -r`. Both the incoming and outgoing connections are removed simultaneously. The PVC must then be removed from each of the network switches involved.

2.5 Debugging

The atmarp utility on the host provides a number of useful options to aid communication debugging. Some of this information can also be displayed for the SCP using various AMI commands. If an equivalent AMI command is available, it will be listed here.

- On the host, atmarp shows the IP address of the entry, the associated ATM address if the mapping is known, the VPI/VCI pair for an established connection, and various flags. The "Classical IP" flag indicates connections using Classical IP over Q.2931. Two other flags are especially important for connection state: "incomplete" indicates that the IP-to-ATM address mapping is not yet known for the given IP address; "pending" indicates that a connection has not (yet) been established. For the switch, AMI command configuration atmarp show provides the same information.
- On the host, atmarp -z displays the ATM address of the given interface. If the host is using ILMI for address registration, the ATM address will be displayed as nonzero when the host has successfully registered its address with the network. For the switch, AMI command configuration atmarp getnsap interface provides the same information.
- On the host, atmarp -g displays the ATM address of the current ARP server for the LIS associated with the given interface. For the switch, AMI command configuration atmarp arpserver show interface provides the same information.
- On the host, atmarp -t indicates whether this host is the ARP server for the LIS associated with the given Classical IP interface. For the switch, AMI command configuration atmarp arpserver show interface provides the same information.
- On the host, atmarp -v displays the status of the ARP server connection. One of the first things to check in a trouble state is that a connection to the ARP server is fully established. If the ARP server connection is established, the ARP VPI/VCI will be nonzero. The CALLING_SERVER flag indicates whether or not an attempt is currently being made to connect to the ARP server. The retry count indicates the number of attempts to connect to the ARP server which have been unsuccessful since the last successful attempt. The InARP count indicates the number of InARP requests that have been sent on the ARP server connection without a response.

2.6 Configuring the Network

In an ATM network, before any connections can be made, the two parties must know each other's NSAP address in order to set up that connection.

To allow those connections to work, the ideal scenario is for all hosts and switches in the network to have support for both ILMI and for RFC-1577 (Classical IP over ATM). However, when using non-Cabletron equipment, this may not be the case. This section will describe how to configure a network with the following scenarios:

- Configuring a third-party host that has no ILMI and no RFC-1577 support
- Configuring a third-party switch that has ILMI support, but no RFC-1577 support
- Configuring a third-party switch that has no ILMI support, but has RFC-1577 support

2.6.1 Third-Party Host with No ILMI and No RFC-1577 Support

To configure a network with a third-party vendor's host (or an edge device) that supports neither ILMI nor RFC-1577 (as shown in Figure 2.1), perform the following steps:

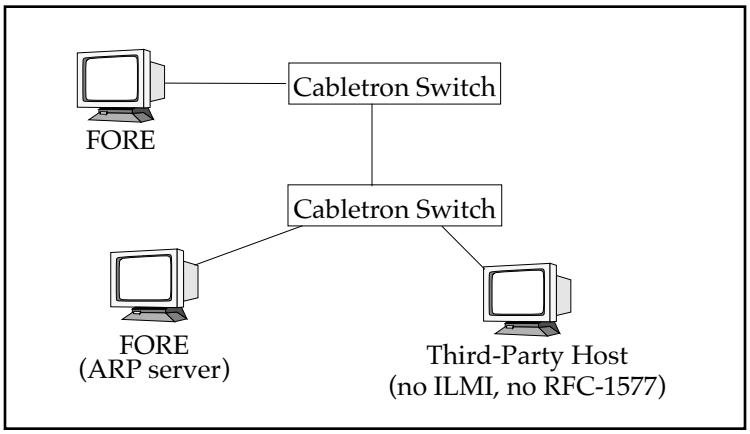

Figure 2.1 - Configuring a Third-Party Host with No ILMI and No RFC-1577 Support

1. Before beginning this process, be sure that FORE software is installed and running on the host that is the ARP server and on all of the other equipment.
2. Using the configuration software of the third-party host, assign that host an NSAP address that has the same prefix as the switch fabric to which it is connected.
3. Configure the switch so that it has a static route to the third-party host using the following AMI command:

configuration nsap route new <NSAP> <mask> -port <port> -vpi <vpi>

Be sure to use a host mask value of 152.

4. Configure the ARP server (edit the firmware download script) with a static IP to NSAP mapping for the third-party host.

2.6.2 Third-Party Switch with ILMI and No RFC-1577 Support

To configure a network with a third-party vendor's switch that supports ILMI, but not RFC-1577, (as shown in Figure 2.2), perform the following steps:

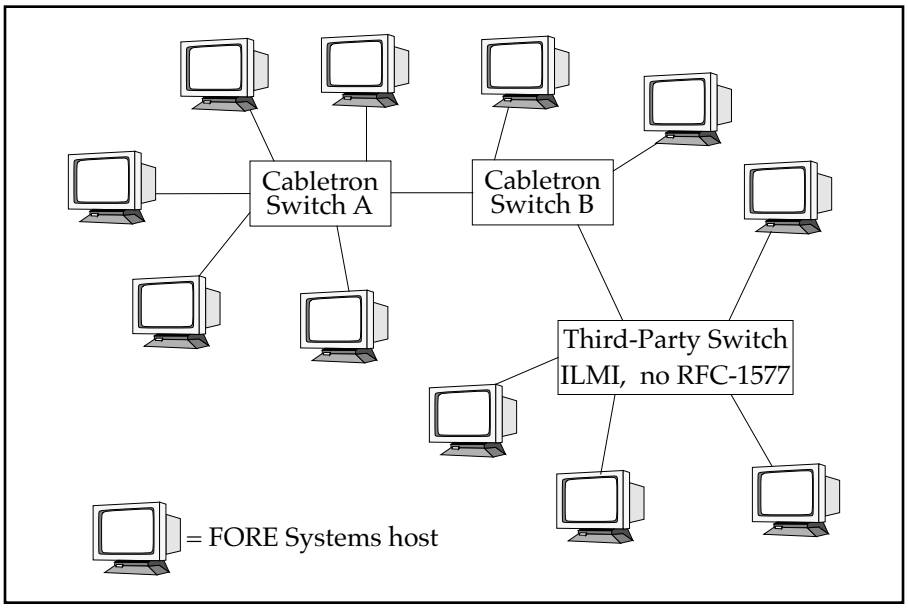

Figure 2.2 - Configuring a Third-Party Switch with ILMI Support and No RFC-1577

1. Be sure that FORE software has been installed on all of the hosts and that ILMI was set in the process. Let ILMI dynamically perform address registration for all of the hosts.
2. Configure a static NSAP route to the third-party switch on Cabletron switch "B" that is physically connected to the third-party switch using the following AMI command:

configuration nsap route new <NSAP> <mask> -port <port> -vpi <vpi>

   Be sure to use a network mask value of 104.

3. Configure two static NSAP routes on the third-party switch, one to each of the Cabletron switches to which the third-party switch is connected, using the third-party vendor's configuration software.

2.6.3 Third-Party Switch with RFC-1577 and No ILMI Support

To configure a network with a third-party vendor's switch that does not support ILMI, but does support RFC-1577 (as shown in Figure 2.3), perform the following steps:

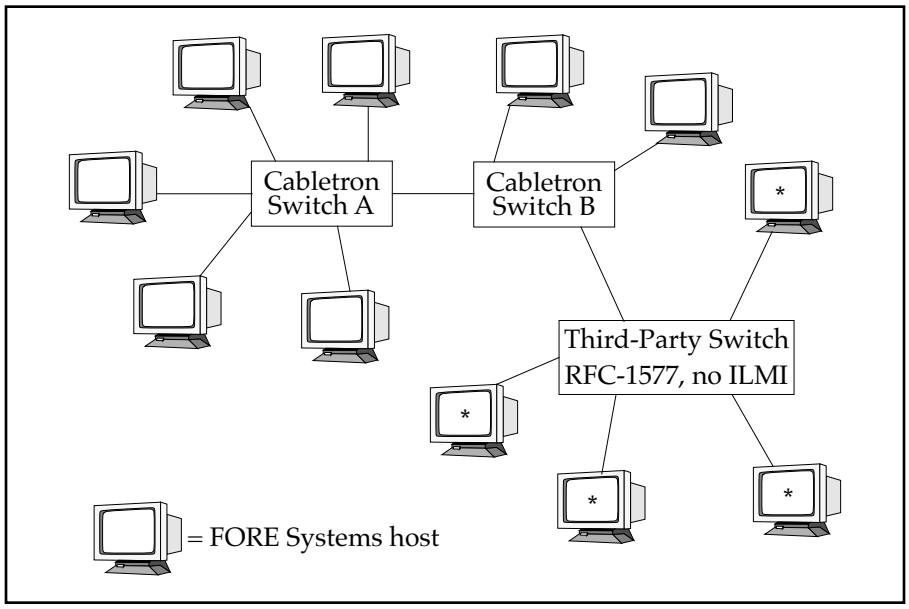

Figure 2.3 - Configuring a Third-Party Switch with RFC-1577 and No ILMI Support

1. Be sure that FORE software has been installed on all of the FORE hosts and that ILMI was set in the process. Let ILMI dynamically perform address registration for all of the FORE hosts and Cabletron switches.
2. Statically configure the non-FORE (*) hosts with NSAP addresses (edit the firmware download script), using the same switch prefix for all of the hosts.
3. Configure a static NSAP route to the third-party switch on Cabletron switch "B" that is physically connected to the third-party switch using the following AMI command:

configuration nsap route new <NSAP> <mask> -port <port> -vpi <vpi>

Be sure to use a network mask value of 104. Also, be sure to use the same prefix that was used to configure the hosts.

4. Configure two static NSAP routes on the third-party switch, one to each of the Cabletron switches to which the third-party switch is connected, using the third-party vendor's configuration software.

CHAPTER 3 Configuring an Emulated LAN

3.1 Introduction

This chapter describes how to design, configure, and maintain an Emulated LAN (ELAN) over an ATM network. An ELAN provides communication of user data frames among all members of the ELAN, similar to a physical LAN. One or more ELANs may run simultaneously (and independently) on the same ATM network. Just as with physical LANs, communication between ELANs is possible only through routers or bridges.

The current software release supports emulation of Ethernet (IEEE 802.3) LANs only. Each ELAN is composed of a set of LAN Emulation Clients (LECs), a LAN Emulation Configuration Server (LECS), a LAN Emulation Server (LES), and a Broadcast and Unknown Server (BUS). Each LEC resides on an ATM host system (PC, Macintosh, UNIX workstation, switch, or bridge device). In the current software release, the LECS, LES, and BUS may reside either in a Cabletron switch or in a UNIX workstation running SunOS version 4.1.x. Additional software features include a colocated BUS (also referred to as an intelligent BUS or a LES/BUS pair) and a LEC "failover" mechanism that can provide redundancy to an ELAN.

3.2 Emulated LAN Components

The components of an emulated LAN include LECs, and LAN Emulation Services consisting of: a LECS, a LES, and a BUS. Each of these services may reside in the same physical system or in separate physical systems. For example, the LECS could reside in a switch, while the LES and BUS reside in a workstation. In the current 4.0 software, the LECS, LES, and BUS are supported only for switches and for SunOS. The functional interconnections of a simple ELAN consisting of two LECs, an LECS, a LES, and a BUS are shown in Figure 3.1.

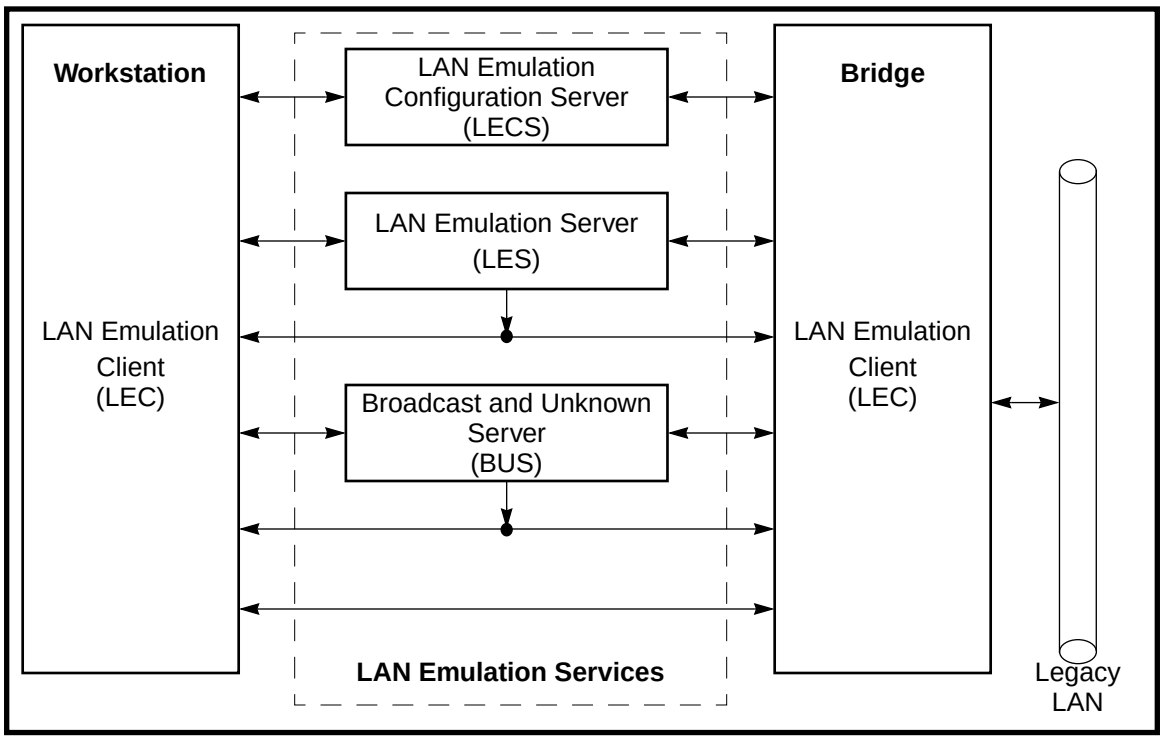

Figure 3.1 - Basic Emulated LAN Interconnections

3.2.1 LAN Emulation Client (LEC)

The LEC is the component in an end system that performs data forwarding, address resolution, and other control functions when communicating with other components within the ELAN. It also provides a MAC level emulated Ethernet interface and appears to higher level software as though a physical Ethernet interface is present. Each LEC must register with both the LES and BUS associated with the ELAN it wishes to join before it may participate in the ELAN. 4.0 supports a maximum of 16 LECs per adapter card.

3.2.2 LAN Emulation Configuration Server (LECS)

The LECS is responsible for the initial configuration of LECs. It provides information about available ELANs that a LEC may join, together with the address of the LES associated with each ELAN. With 4.0, the user may also use the LECS to associate multiple LESs with a given ELAN. This feature allows LECs to "failover" to a hierarchy of redundant services if the primary LES for an ELAN goes down.

3.2.3 LAN Emulation Server (LES)

The LES implements the control coordination function for the ELAN. The LES provides the service of registering and resolving MAC addresses to ATM addresses. A LEC registers its own address with the LES. A LEC also queries the LES when the client wishes to resolve a MAC address to an ATM address. The LES either responds directly to the client or forwards the query to other clients so they may respond. There is only one instance of an active LES per ELAN.

3.2.4 Broadcast and Unknown Server (BUS)

Unlike traditional shared-media LAN architectures such as Ethernet, ATM is connection based. Therefore, it has no built-in mechanism for handling connectionless traffic such as broadcasts, multicasts, and unknown unicasts. In an emulated LAN, the BUS is responsible for servicing these traffic types by accepting broadcast, multicast, and unknown unicast packets from the LECs to the broadcast MAC address (FFFFFFFFFFFF) via dedicated point-to-point connections, and forwarding the packets to all of the members of the ELAN using a single point-to-multipoint connection. Each LEC is associated with only one active BUS. 4.0 also supports the use of a colocated BUS (also referred to as an intelligent BUS or a LES/BUS pair) that allows the BUS to use the LES's registration table to direct unknown unicast traffic.

3.3 Emulated LAN Operation

This section describes the operation of an ELAN and its components from the point of view of a LEC. The operation of an ELAN may be divided into three phases:

1. Initialization
2. Registration and Address Resolution
3. Data Transfer

ELAN components communicate with each other using ATM connections. LECs maintain separate connections for traffic control functions and data transfer. The following connection types are used by the LEC when operating in an ELAN:

- *Configuration-Direct Connection*: a bidirectional point-to-point VCC set up by the LEC to the LECS.
- *Control-Direct Connection:* a bidirectional point-to-point VCC set up by the LEC to the LES. This connection must be maintained for the duration of the LEC's participation in the ELAN.
- *Control-Distribute Connection*: a unidirectional point-to-multipoint VCC set up by the LES to the LEC. This connection must be maintained for the duration of the LEC's participation in the ELAN.
- *Multicast-Send Connection*: a bidirectional point-to-point VCC set up by the LEC to the BUS for sending multicast data to the BUS. The LEC must attempt to maintain this connection while participating in the ELAN.
- *Multicast-Forward Connection*: a unidirectional point-to-multipoint VCC set up from the BUS to LECs participating in the ELAN. This VCC must be established before a LEC participates in an ELAN. The LEC must attempt to maintain this connection while participating in the ELAN.
- *Data-Direct Connection*: a bidirectional point-to-point VCC set up between LECs that want to exchange unicast data traffic.

For the following discussion, please refer to Figure 3.2.

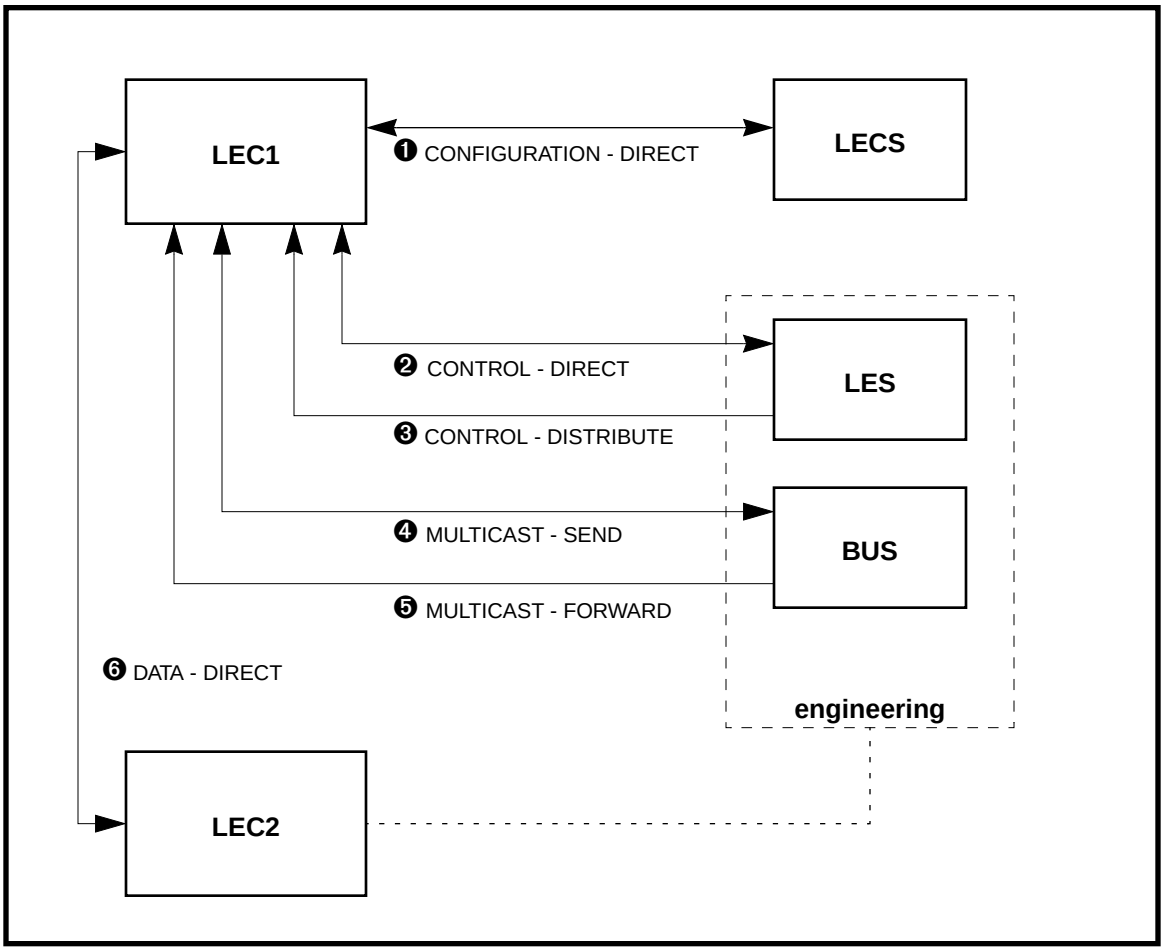

Figure 3.2 - ELAN Operation

3.3.1 Initialization

Upon initialization, LEC1 obtains its own ATM address via address registration. LEC1 obtains the address of the LECS in one of three ways: via manual configuration, via the "well-known" address defined by the ATM Forum's LANE standards (47.0079.00.000000.0000.0000.0000.00A03E000001.00), or via PVC (0,17).

Once it knows the location of the LECS, LEC1 establishes a configuration-direct connection ❶ to the LECS. When connected, the LECS provides LEC1 with the information necessary to connect to the ELAN it wishes to join. This information includes such parameters as: the ATM address of the ELAN's LES, the type of LAN being emulated, the maximum packet size, and the name of the ELAN (`engineering`, for example). This configuration information is contained in a configuration file that must be built and maintained by the network administrator.

> **NOTE:** Detailed information about the LECS configuration file may be found in Section 3.4.1.

3.3.2 Registration and Address Resolution

After obtaining the address of the LES, LEC1 establishes a control-direct connection ❷ to the LES. The LES assigns LEC1 a unique identifier, and LEC1 registers its own MAC and ATM addresses with the LES. (The LES maintains a table containing the MAC addresses and corresponding ATM addresses of all members of the ELAN.) At this point, LEC1 has "joined" the ELAN.

The LES then establishes a control-distribute connection ❸ back to LEC1. Connections ❷ and ❸ can now be used by LEC1 to send LAN Emulation ARP (LE_ARP) requests to the LES, and receive replies.

LEC1 now sends an LE_ARP request to the LES to get the ATM address of the BUS corresponding to the broadcast MAC address (FFFFFFFFFFFF). The LEC then establishes a multicast-send connection ❹ to the BUS. The BUS responds by setting up a multicast-forward connection ❺ to the LEC.

At this point, the LEC is ready to transfer data.

3.3.3 Data Transfer

When LEC1 receives a network-layer packet from a higher layer protocol to transmit to some destination MAC address (for example, LEC2), LEC1 initially does not know the corresponding ATM address of the destination. Consequently, LEC1 transmits an LE_ARP request to the LES.

> **NOTE:** The example shown in Figure 3.2 assumes that LEC2 has already registered with the LES, and that connections similar to those described for LEC1 already exist.

While waiting for the LES to respond, LEC1 forwards the packet to the BUS. The BUS broadcasts the packet to all LECs on the ELAN. This is done to avoid data loss, and to minimize connection set-up latency (due to the LE_ARP process) that may not be acceptable to some network protocols.

If the LE_ARP response is received, LEC1 establishes a data-direct connection ❻ to the destination address of LEC2. This path will be used for subsequent data transfers. Before LEC1 begins to use this connection, it first sends a "flush" packet via the BUS to the destination, LEC2. When LEC2 acknowledges receipt of this packet, signifying that the BUS path is empty, only then does LEC1 begin to use the data-direct connection ❻ for data transfer. This process ensures that the network protocol's frames arrive in the proper order.

If no response is received to the LE_ARP, LEC1 continues to send data via the BUS, while continuing to LE_ARP until a response is received and a data-direct connection to LEC2 is established.

If LEC1 already has a data-direct connection to a MAC address it wishes to reach, it need not go through the LE_ARP process again. Instead, it continues to use the current connection. This is possible because each LEC maintains a cache of MAC address to ATM address mappings that it receives in response to the LE_ARPs it has sent. Entries in this cache are "aged" out over a period of time. Data-direct connections are also cleared if they remain inactive for a period of time.

3.4 Configuring an ELAN

To configure an ELAN on a switch, you must log into AMI on a switch running software version 4.0 and use the commands found under `configuration lane`.

> **NOTE:** More information about each of these commands may be found in Appendix B of this manual.

There are three major steps that the system administrator should follow in order to configure and maintain ELANs:

1. Configure an LECS configuration database file.
2. Start the LAN Emulation Services (LECS, LES, and BUS).
3. Start the LEC(s) and join an ELAN.

> **NOTE:** Steps 1 may be performed using a text editor on any system. However, for software version 4.0, the resulting file can be **used** only on systems running under SunOS Version 4.1.x or on a switch.

The remainder of this section gives a practical example of configuring and administering an ELAN using software version 4.0.

3.4.1 Configuring an LECS Configuration Database File

The LECS uses a text configuration file to contain the configuration information needed by LECs that wish to participate in an ELAN. The LECS configuration file may be built and edited using a text editor such as vi or emacs.

3.4.1.1 Before You Begin

Before building or modifying the LECS configuration file, you should first determine the topology of the ELAN or ELANs that you want to administer. You must supply the following information when building or editing the LECS configuration file:

- Provide the name of each ELAN (engineering, marketing, etc.).
- Provide the ATM address of the LES for each ELAN.
- If LEC failover is to be made available to clients on a given ELAN, the ATM addresses for all LESs in that ELAN's failover sequence must be included.
- Provide the address of each LEC that may participate in each ELAN.
- If you wish LECs to use a default ELAN, the default LES information must also be included.
- Provide various other configurable parameters.

CAUTION

Do not attempt to edit an existing functional LECS configuration file without first making a backup copy of the file. Incorrect modification of the configuration file could result in loss of communication between one or more members of a defined ELAN. It may also result in one or more of the ELAN(s) defined in the file going "down".

NOTE: You may make changes to the LECS configuration file while the LECS process is running. The configuration file is reread periodically by the LECS process (the default period is ten minutes). Consequently, any changes that you make to the configuration file are not recognized until the file is reread.

3.4.1.2 LECS Configuration File Syntax

Each line that you enter in the configuration file takes the general form:

```
[[group].]key : value
```

The group field may represent:

- ELANs (by name) - ELAN names are case-sensitive, and may not exceed 32 characters in length
- clients ATM or MAC addresses
- miscellaneous LECS control information specified by using a group name of LECS

The key field is used to denote an individual parameter within a group.

The value field contains the value assigned to the key.

Omitting the group implies that the key and value apply to all groups in the configuration file. Leading and trailing spaces, as well as spaces on either side of the ":", are ignored.

For example, to specify a maximum frame size of 1516 bytes for the ELAN named engineering, enter the following:

```
engineering.Maximum_Frame_Size : 1516
```

Similarly, to specify a default maximum frame size of 1516 bytes for all ELANs defined in a given configuration file, enter the following:

```
.Maximum_Frame_Size : 1516
```

Table 3.1 defines the various key parameters that may be entered in the configuration file. The default value for each parameter is also given.

Table 3.1 - LECS Configuration File Parameters

Parameter	Definition
`.LAN_Type: Ethernet/IEEE 802.3`	Identifies the type of emulated LAN, either Ethernet/IEEE 802.3 or Token Ring.
`.Maximum_Frame_Size: 1516`	The length (in number of bytes) of the largest frame field. Selections are: 1516, 4544, 9234, and 18190. Currently, only the default value of 1516 is supported.
`.Control_TimeOut: 120`	Specifies the timing out of request/response control frame interactions, in seconds. Default is 120 seconds.
`.Maximum_Unknown_Frame_Count: 1`	Limits the number of unicast frames sent to the BUS. Default is 1 frame.
`.Maximum_Unknown_Frame_Time: 1`	Limits the number of unicast frames sent to the BUS in the specified number of seconds. Default is 1 second.
`.VCC_TimeOut_Period: 1200`	Specifies the length of time that an idle data connection remains open before being closed. Default value is 1200 seconds.
`.Maximum_Retry_Count: 1`	Limits the number of LE_ARP requests. Default is 1.
`.Aging_Time: 300`	Specifies the period that LE_ARP cache table entries remain valid, in seconds. Default value is 300 seconds.
`.Forward_Delay_Time: 15`	Specifies the timing out of non-local ARP cache entries in seconds. Default value is 15 seconds.
`.Expected_LE_ARP_Response_Time: 1`	Specifies the maximum time a LEC expects an LE_ARP request/response will take, in seconds. Default value is 1 second.
`.Flush_TimeOut: 4`	Specifies the maximum time a LEC expects an LE_FLUSH request/response will take, in seconds. Default value is 4 seconds.
`.Path_Switching_Delay: 6`	Minimum time between switching BUS and data paths, in seconds. Default value is 6 seconds.
`.Multicast_Send_VCC: Best Effort`	Specifies the multicast send mode, either Best Effort, Variable, or Constant.
`.Connection_Complete_Timer: 4`	Specifies the time period in which data or READY_IND is expected, in seconds. The default is 4 seconds.

Lines beginning with # may be inserted if you wish to include comments or to improve the clarity of the presentation when the file is viewed or printed. These lines are ignored when the file is read. Lines may be continued by escaping the end-of-line with a backslash "\" (do not enter the quote marks).

3.4.1.3 Defining an ELAN

Each ELAN is defined by an address statement whose value denotes the ATM address of the ELAN's LES. For example:

```
engineering.Address: 47000580ffe1000000f21a01b90020480605b211
```

If you wish to define multiple instances of an ELAN for use by the LEC failover mechanism, the address statements would appear as follows:

```
marketing|0.Address: 47000580ffe1000000f21a01b90020480605b221
marketing|1.Address: 47000580ffe1000000f21a01b90020480605b223
marketing|2.Address: 47000580ffe1000000f21a01b90020480605b225
```

In this example, redundant LES addresses for the marketing ELAN have been defined. Note that the group name, marketing, ends with the meta-character |, followed by a number indicating the position of the LES in the failover sequence. If the marketing|0 LES were to fail, LECs could "failover" to the marketing|1 LES. The LEC then periodically tries to reconnect with marketing|0. If marketing|1 also failed, the LEC would then failover to marketing|2, (and periodically attempt to reconnect to marketing|0 and marketing|1).

In addition, you may instruct a given ELAN to override any of the default values. For example, the engineering ELAN could override the default Maximum_Frame_Size of 1516; thus:

```
engineering.Maximum_Frame_Size: 4544
```

If you want to control which clients may or may not join a given ELAN, two additional keys, `Accept` and `Reject`, whose values are comma-separated lists of matching elements, may be used. These values may be:

a MAC address,

```
engineering.Accept: 0020480605b2 , 002048080011 , 0020481020ef
```

an ATM address and equal-length bit mask,

```
engineering.Accept: 47000580FFE1000000000000000000204800000000 \
                        FFFFFFFFFFFFFF000000000000FFFFFF00000000
```

or an ATM address containing “don’t-care” semi-octets denoted by an “X”.

```
marketing.Accept: 47000580FFE100XXXXXXXXXXXX002048XXXXXXXX
```

The last two forms of ATM-address matching elements are functionally the same. The latter is shorter but only allows for masks whose semi-octets are all ones or all zeros, while the former allows for arbitrary masks. A prospective-client address is “captured” by an ELAN name if the client’s address matches one of the `Accept` elements but not one of the `Reject` elements (if present). Finally, an ELAN may be configured to accept any client that wishes to join by including the following statement:

```
default.Accept:XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

The order in which to apply the Accept and Reject rules is given by a `Match.Ordering` group.key statement, whose value is a comma-separated list of ELAN names. For example:

```
Match.Ordering: default, engineering, marketing|0, marketing|1,\
                        marketing|2
```

The names of all ELANs that have `Accept` keys must be included in `Match.Ordering`. Names of all of the members of a LEC failover sequence must be included so that LECs may use them.

3.4.1.4 Defining a Client

Clients need not be defined in the LECS configuration file. Typically, you would define a client for the purpose of overriding one or more of the default configuration parameters for that particular client.

A client is defined by using its ATM or MAC address in the group field, and perhaps giving the name of its ELAN as the value of the LAN_Name key. For example:

```
47000580FFE1000000000000000204822222222.LAN_Name: engineering
002048ABCDEF.LAN_Name: marketing
```

Configuration parameter overrides can also be given on a per-client basis. For example, the following statements override the default VCC_TimeOut_Period and Aging_Time configuration parameters for a client whose MAC address is 002048080011 on the engineering ELAN:

```
002048080011.LAN_Name:engineering
002048080011.VCC_TimeOut_Period:1200
002048080011.Aging_Time: 30
```

3.4.1.5 LECS Control Parameters

Specifying values for keys in the LECS group provides control over the operation of the LECS process.

> **NOTE:** If you change the values of the LECS control parameters while the LECS process is running, the new values do not take effect until the LECS process is stopped, and then restarted.

When a client contacts the LECS, the connections established are known as Configuration Direct VCCs. To override the default value of the VCC_TimeOut_Period key (the number of seconds before an idle Configuration Direct VCC is automatically closed by the LECS), enter a statement similar to the following:

```
LECS.VCC_TimeOut_Period: 1200
```

The LECS periodically checks whether its configuration file has been modified, and, if it has, the file is reread. The length of this period, in seconds, is given by the Reload_Period key:

```
LECS.Reload_Period: 600
```

The Permanent_Circuits key holds a comma-separated list of VPI.VCI pairs denoting the local ends of 0.17 PVCs on which the LECS should listen. For example:

```
LECS.Permanent_Circuits: 0.42, 0.112
```

The LECS can provide the client with a fourteen-bit pattern to permute the MAC-address generation algorithm. This bit pattern is specified with the MAC_Address_Base key.

```
LECS.MAC_Address_Base: 38fe
```

3.4.2 Sample LECS Configuration File

CAUTION

Do not attempt to edit an existing functional LECS configuration file without first making a backup copy of the file. Incorrect modification of the configuration file could result in loss of communication between one or more members of a defined ELAN. It may also result in one or more of the ELAN(s) defined in the file going "down".

NOTE: For a detailed discussion of how to configure an LECS configuration file similar to the one given in this section, please refer to Section 3.4.1.

The sample LECS configuration file shown at the end of this section in Figures 3.3 and 3.4 defines three ELANs:

- default
- engineering
- marketing

The Match.Ordering statement specifies the ELAN names in the order that prospective clients will attempt to join. Note that redundant marketing ELANs have been included. This allows clients attempting to join the marketing ELAN to "failover" to one of the redundant LES instances if the primary LES fails.

The default configuration parameters are shown with their default values. These values apply to all ELANs in this configuration file, unless overridden for a particular ELAN or client.

ELAN default is configured to accept any client that wishes to join. The administrator must substitute the ATM address of the default LES in place of the "4"s in the default.Address statement.

ELAN engineering has overridden the default Maximum_Frame_Size with a new size of 4544 bytes. Consequently, this frame size applies only to traffic on the engineering ELAN. The default and marketing ELANs continue to use the default frame size of 1516 bytes.

ELAN marketing has three instances, |0, |1, and |2. These three instances provide LECs that join the ELAN marketing the information necessary to implement the LEC "failover" mechanism. marketing|0 defines the primary services for the ELAN. If marketing|0 fails, LECs connected to it can automatically switch to marketing|1 and continue operating (while periodically attempting to rejoin marketing|0). If marketing|1 were also to fail, the LECs could then switch to services on marketing|2.

Two LECs, whose MAC addresses are 002048080011 and 0020481020ef, are identified as acceptable clients for the engineering and marketing ELANs.

```
#
# The search ordering of elan names
#
Match.Ordering: default, engineering, marketing|0, marketing|1, marketing|2
#
# the default configuration parameters
#
.Control_TimeOut: 120
.Maximum_Unknown_Frame_Count: 1
.Maximum_Unknown_Frame_Time: 1
.VCC_TimeOut_Period: 1200
.Maximum_Retry_Count: 1
.Aging_Time:  300
.Forward_Delay_Time: 15
.Expected_LE_ARP_Response_Time: 1
.Flush_TimeOut: 4
.Path_Switching_Delay: 6
.Multicast_Send_VCC_Type: Best Effort
.Connection_Complete_Timer: 4
.LAN_Type:    Ethernet/IEEE 802.3
.Maximum_Frame_Size: 1516
#
# Parameters for the active default elan
#
default.Address:44444444444444444444444444444444444444
default.Accept: XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

Figure 3.3 - Sample LECS Configuration File (Part One of Two)

```
#
# Parameters for elan: engineering
#
engineering.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.11
engineering.Accept: 002048080011 , 0020481020ef
engineering.Maximum_Frame_Size: 4544
#
# Parameters for elan: marketing|0
#
marketing|0.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.21
marketing|0.Accept:  002048080011 , 0020481020ef
#
# Parameters for elan: marketing|1
#
marketing|1.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.23
marketing|1.Accept:  002048080011 , 0020481020ef
#
# Parameters for elan: marketing|2
#
marketing|2.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.25
marketing|2.Accept:  002048080011 , 0020481020ef
```

Figure 3.4 - Sample LECS Configuration File (Part Two of Two)

3.4.3 The Default LECS Configuration File

A default LECS configuration file is installed with software version 4.0. It is reproduced below for reference.

```
#
# This sample LECS configuration file contains the minimum information
# needed to start a default ELAN that all clients can join.  You MUST
# replace the '4's below with the ATM address of the LES.
#

#
# The search ordering of ELAN names
#
Match.Ordering: default

#
# Parameters for the default ELAN
#
default.Address: 44444444444444444444444444444444444444
default.Accept:  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

Figure 3.5 - Default LECS Configuration File

NOTE: Make a backup copy of the supplied LECS configuration file and work with the copy.

As supplied, this file allows you to set up a default ELAN that accepts any client that wants to join. Before you can use the supplied file, you must modify it as follows:

1. Obtain the ATM address of the machine where the LES for the default ELAN will be started. You must then modify the machine's ATM address by changing the selector byte. This modified address is the same address that is used when starting the default LES.
2. Substitute the address determined in step 1 for the string of "4"s in the default.Address statement of the LECS configuration file. The default.Accept string of "X"s should not be changed as this parameter allows any client who wishes to join the default ELAN.
3. Save the modified file.

NOTE: For a detailed discussion of how to configure an LECS configuration file other than the supplied default file, please refer to Section 3.4.1.

3.4.4 Starting the LAN Emulation Services

LAN Emulation services include the LECS, LES, and BUS. Once the LECS configuration database file has been configured, each of these services must be started so that they are available for LECs to attempt to use. These services may, but need not, run in the same machine.

3.4.4.1 Starting the LECS

Once an LECS configuration file has been configured, you need to retrieve the LECS configuration database file that you built by using the following AMI command:

configuration lane lecs get <host>:<remotefile> [<localfile>]

For example, you would enter something similar to the following:

configuration lane lecs get 198.29.22.46:lecs.cfg

NOTE: The default local file for a 9A000, an SFCS-200BX, an SFCS-200WG, and an SFCS-1000 is `lecs.cfg`.

NOTE: On a 9A000, an SFCS-200BX, an SFCS-200WG, and an SFCS-1000, this file is retrieved via tftp.

After you have retrieved the LECS configuration database file, use the following AMI command to start the LECS service on the switch:

configuration lane lecs new <LECS Selector byte (HEX)> [-db <LECS database file>]
[-default <LES atm address>]

For example, to start the LECS service using only the `-db` option, you would

enter something similar to the following:

configuration lane lecs new 0x0c -db lecs.cfg

Use the following AMI command to verify that the LECS has been started and is running. The OperStatus field should display up, meaning that the LECS is enabled. Because the -default option has not been specified, the Default LES field is all zeros. If you use the -default option, then that LES address would be displayed in the Default LES field.

configuration lane lecs show

```
Index  AdminStatus  OperStatus   Selector  Database
    1  up           up               0x0c  lecs.cfg
       Default LES :00000000000000000000000000000000000000000000
```

3.4.4.2 Creating a LES and a BUS

Next, the LES and BUS services must be started for the ELAN. The preferred method is to use the `colocated_bus` option when creating the LES, which means that the LES and BUS services for a particular ELAN will be started together and will be running on the same switch. This configuration may provide better BUS performance. To do this, use the following AMI command:

**conf lane les new <LES Selector byte (HEX)> <BUS ATM address> <LES name> **

[(colocated_bus)]

For example, you would enter something similar to the following:

conf lane les new a b default colocated_bus

NOTE: By using the `colocated_bus` option, you are creating a LES and BUS using a single AMI command. There is no need to create a BUS separately.

Use the following AMI command to verify that the LES and the BUS have been started and are running. The `OperStatus` field in this listing should display `up`, meaning that the LES and the BUS are enabled.

configuration lane les show

```
Index  AdminStatus  OperStatus   Selector  ELAN
    1  up           up              0x0a  default
      BUS:0x47.0005.80.ffe100.0000.f215.096b.00204815096b.0b (Co-Located)
```

3.4.5 Starting the LEC(s) and Joining an ELAN

Now that the ELAN services have been started, you can have LECs join the ELAN that you have created.

NOTE: The switch software only allows you to create an instance of a LEC on a switch. To create an instance of a LEC on a host, you must use the *ForeRunner* VLAN Manager or use a *ForeRunner* host adapter. Please refer to the respective User's Manual for instructions.

3.4.5.1 Creating a LEC

To start a LEC that will attempt to join the ELAN, use the following AMI command:

conf lane lec new <LEC Selector byte (HEX)> <ELAN name> [(automatic | manual)]

manual mode options: [-lecs <LECS address>] or [-les <LES address>]

NOTE: The recommended method for starting a LEC is to use the `automatic` mode, meaning that the LEC will attempt to contact the LECS on the "well-known" address as defined by the ATM Forum's LAN Emulation standards (47.0079.00.000000.0000.0000.0000.00A03 E000001.00).

For example, to start a LEC that attempts to join the default ELAN called `default` (assuming that a default ELAN has been defined in the LECS configuration file), enter the following:

configuration lane lec new 0 default

If you decide to use the manual mode, you must enter either a LECS address other than the well-known address or you must enter a LES address. If you enter a LES address, this means that the LEC bypasses the LECS and directly contacts the specified LES.

For example, you would enter something similar to the following:

configuration lane lec new aa default manual -les

47000580ffe1000000f21a00d50020481a00d50c

Regardless of which way you start the LEC, you can verify that the LEC has joined the ELAN by using the following AMI command:

configuration lane lec show

```
       Admin   Oper
Index  Status  Status  Sel   Mode        MACaddress    IfName    ELAN
    1  up      up       0x00  automatic   000000000000   el0    default
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.a0
```

Because the OperStatus field is displaying joining, this means that the LEC is still registering with the ELAN. When it has finished, the OperStatus field displays up.

After the first LEC has joined the ELAN, you can use the VLAN Manager or the host software to add more LECs to this ELAN. Once all the LECs have joined, the ELAN is complete.

3.4.5.2 Configuring the LEC Failover Mechanism

An added feature of 4.0 software is the LEC failover mechanism. This feature allows multiple redundant services to be configured for an ELAN as shown in the sample LECS configuration file in 3.4.2. Consequently, to start LECs that join ELAN marketing and make use of the failover mechanism set up in the sample LECS configuration file, enter the following AMI commands:

```
configuration lane lec new aa marketing|0
```

followed by,

```
configuration lane lec new ab marketing|1
```

and then,

```
configuration lane lec new ac marketing|2
```

The above commands start three LECs. Initially, the LEC that connects to ELAN marketing|0 services is the active LEC. If marketing|0 services fail, the LEC associated with ELAN marketing|1 services becomes active. Periodically, the LEC associated with ELAN marketing|0 attempts to re-establish service. If both marketing|0 and marketing|1 services fail, the LEC associated with ELAN marketing|2 then becomes active. You may define as much redundancy into your system as resources allow.

CAUTION

Deleting (stopping) the currently active failover LEC removes the network interface (e.g., el0) associated with **all** of the failover LECs, causing the other LECs in the failover mechanism to lose IP connectivity.

CHAPTER 4 SONET Configuration

SONET network modules may require some additional configuration. There are several parameters that can be configured in order to work with other SONET equipment or to perform testing on the SONET ports.

4.1 SONET Front Panel LEDs

There is an LED corresponding to both the transmit and receive lines of each SONET port. These LEDs provide some information about the state of the port depending on their color.

4.1.1 Transmit Indicators

The LED corresponding to the transmit line of the port has the following meanings depending on its color:

off Indicates that there are currently no cells being transmitted from the port.

green Indicates that a cell is being transmitted on the port. The intensity of the green light increases as the traffic on the port increases.

yellow Indicates that a yellow alarm is being received from the far end of the connection. A yellow alarm indicates that the transmit signal is not being received by the far end.

NOTE: There is no red LED on the transmit indicators. However, a yellow LED may sometimes appear to be red.

4.1.2 Receive Indicators

The LED corresponding to the receive line of the port has the following meanings depending on the color:

off Indicates that a carrier has been detected on the line. A carrier is detected when there is a proper light signal on the line.

green Indicates that a cell is being received on the port. The intensity of the green light increases as the traffic on the port increases.

red Indicates a loss of carrier.

4.2 Configuring SONET Mode

There are two modes of operation that the SONET ports support: `sonet` and `sdh`. These modes of operation affect the significance of the header bits in the SONET frames. The default mode of operation is `sonet`.

CAUTION

During standard system operation, `sonet` mode transmits unassigned cells and `sdh` mode transmits idle cells. If the type of operation is changed from `sonet` to `sdh`, the type of cells being transmitted does not automatically change. The user must be aware if they wish to be sending unassigned or idle cells for both `sonet` and `sdh` operation. If the type of cell currently selected is not the type the user wants to send, they can change the cell type being sent via the ATM Management Interface (AMI).

To change the mode of a SONET port, log in to AMI. (Please refer to Appendix A of this manual for information about logging into AMI). Enter the following parameters:

configuration port sonet mode <port> (sonet|sdh)

The <port> variable indicates the SONET port that is to be modified and the (sonet|sdh) variable indicates the mode of operation that is to be used for the specified port. If there are no SONET ports on the switch fabric, then this option is disabled.

After the mode of operation is modified on a SONET port, the switch control processor (SCP) immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

4.3 Configuring SONET Empty Cells

To change the type of cells sent as empty cells (filler that is sent when a port is not sending data) on a SONET network module port, log in to AMI and enter the following parameters:

configuration port sonet emptycells <port> (idle | unassigned)

The <port> variable indicates the SONET port that is to be modified and the (idle|unassigned) variable indicates the type of cells that the specified port sends as filler when the port is not sending data. The default setting is unassigned. If there are no SONET ports on the switch fabric, then this option is disabled.

After the type of empty cells is modified on a SONET port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

4.4 Configuring SONET Loopback

To facilitate testing of the SONET ports, there are two different loopback configurations available: line and diagnostic. When a SONET port is in loopback mode, it no longer passes normal traffic.

To change the loopback state on a SONET port, log in to AMI and enter the following parameters:

configuration port sonet loopback <port> (line | diag | none)

The <port> variable indicates the SONET port that is to be modified and the (line | diag | none) variable indicates the type of loopback to be used on the specified port. The default loopback setting is none which means that no loopback will take place on that port. If there are no SONET ports on the switch fabric, then this option is disabled.

After the loopback mode is modified on a SONET port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

4.4.1 Diagnostic Loopback

Diagnostic loopback connects the receiver to the transmitter. The SONET stream being transmitted by the SCP to a port is looped back to the SCP. The stream is still transmitted over the fiber, but the incoming stream is ignored.

4.4.2 Line Loopback

Line loopback connects the transmitter to the receiver. The data stream received from the fiber is retransmitted back out to the fiber. In line loopback, the port acts as if it were an optical repeater. Cells generated by the SCP to this port are not sent over the fiber.

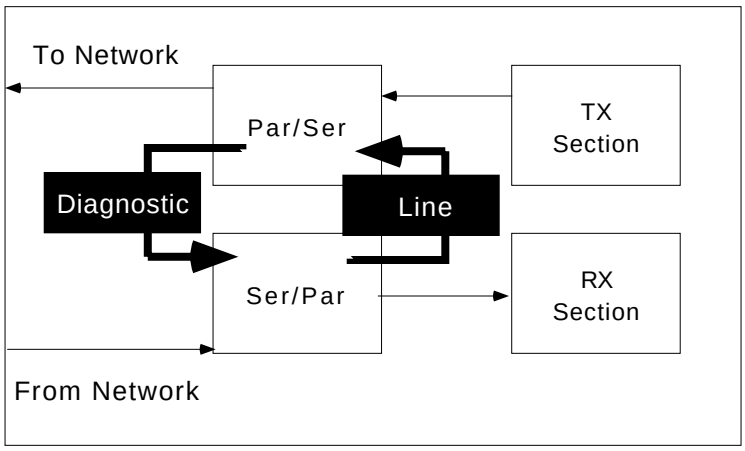

Figure 4.1 - SONET Single Port Loopback Diagram

4.5 Displaying SONET Error Counters

AMI allows the user to display several SONET counters. The counters can be accessed by logging in to AMI. Enter the following parameters at the prompt to display the SONET network module statistics:

statistics sonet

```
sonet Port 1A1 Counter              Value           Delta
------------------------------  --------------  --------------
sonetSectionBIPs                   383833630         1584162
sonetSectionLOSs                       22103              99
sonetSectionLOFs                       23991              99
sonetLineBIPs                            532               0
sonetLineFEBEs                           355               0
sonetLineAISs                          23991              99
sonetLineFERFs                             6               0
sonetPathBIPs                            244               0
sonetPathFEBEs                           211               0
sonetPathLOPs                           1888               0
sonetPathAISs                          23991              99
sonetPathYellows                       23997              99
sonetAtmCorrectableHCSs                    0               0
sonetAtmUncorrectableHCSs                  0               0

Press return for more, q to quit: q
```

4.6 SONET Error Counter Descriptions

sonetSectionBIPs The number of Section BIP-8 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. The calculated BIP-8 code is compared with the BIP-8 code extracted from the B1 byte of the following frame. Differences indicate that a section level bit error has occurred.

sonetSectionLOSs The number of seconds in which Loss Of Signal (LOS) has occurred. A LOS is declared when a pattern of all zeros is detected for 20 +/- 3ms. LOS is cleared when two valid framing words are detected and during the intervening time no LOS condition is detected.

sonetSectionLOFs The number of seconds in which Loss Of Frame (LOF) has occurred. A LOF is declared when an out-of-frame (OOF) condition persists for 3ms. The LOF is cleared when an in-frame condition persists for 3ms. While in-frame, the framing bytes (A1, A2) in each frame are compared against the expected pattern. Out of frame is declared when four consecutive frames containing one or more framing pattern errors have been received.

sonetLineBIPs The number of Line BIP-24 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. The calculated BIP-24 code is based on the line overhead and synchronous payload envelope (SPE) of the STS-3c stream. The line BIP-24 code is a bit interleaved parity calculation using even parity. The calculated code is compared with the BIP-24 code extracted from the B2 bytes of the following frame. Differences indicate that a line layer bit error has occurred.

sonetLineFEBEs The number of line Far End Block Errors (FEBE) that have been detected since the last time the port has been reset.

sonetLineAISs The number of seconds in which line Alarm Indication Signal (AIS) has occurred. A line AIS is asserted when a 111 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line AIS is removed when any pattern other than 111 is detected in these bits for five consecutive frames.

sonetLineFERFs The number of seconds in which line Far End Receive Failure (FERF) has occurred. A line FERF is asserted when a 110 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line FERF is removed when any pattern other than 110 is detected in these bits for five consecutive frames.

sonetPathBIPs The number of Path BIP-8 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. A path BIP-8 error is detected by comparing the path BIP-8 byte (B3) extracted from the current frame, to the path BIP-8 computed for the previous frame.

sonetPathFEBEs The number of path Far End Block Errors (FEBE) that have been detected since the last time the port has been reset. FEBEs are detected by extracting the 4-bit FEBE field from the path status byte (G1). The legal range of the 4-bit field is between 0000 and 1000, representing zero to eight errors. Any other value is interpreted as zero errors.

sonetPathLOPs The number of seconds in which path Loss Of Pointer (LOP) has occurred. A path LOP is detected when a 'normal pointer value' is not found in eight consecutive frames. The LOP is cleared when a 'normal pointer value' is detected for three consecutive frames.

sonetPathAISs The number of seconds in which path Alarm Indication Signal (AIS) has occurred. A path AIS is asserted when an all-ones pattern is detected in the pointer bytes (H1 and H2) for three consecutive frames. It is cleared when a valid pointer is detected for three consecutive frames. AIS indicates an upstream failure has been detected.

sonetPathYellows The number of seconds in which path yellow alarm has occurred. A path yellow alarm is detected by extracting bit 5 of the path status byte. If bit 5 is high for ten consecutive frames, a yellow alarm is declared. A yellow alarm is cleared when bit 5 is low for ten consecutive frames. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

sonetAtmCorrectable-HCSs The number of correctable Header Check Sequence (HCS) error events that occurred since the port was reset. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

sonetAtmUncorrect-ableHCSs The number of uncorrectable Header Check Sequence (HCS) error events that occurred since the port was reset. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

4.7 Configuring SONET Timing

The SONET ports on an individual switch fabric can derive timing from one of two sources: an internal clock (internal) or the incoming SONET data stream (network). The default setting is internal clocking. Cabletron Systems recommends that the following clocking conventions be used to configure the SONET clocking on an individual switch fabric:

- When connecting to a carrier-provided SONET service, the carrier's recommendation should be followed. In most cases, the carrier provides a clock and the SCP should be configured for network timing.
- When connecting two switch fabrics over a continuous fiber within the campus, both ends of the connection should be set to internal clocking (default).
- When connecting two switch fabrics through other SONET equipment within the campus, it may be necessary to change the clocking to network (depending on the type of equipment). In this case, the recommendation of the campus network administration should be followed.

To change the clocking source for a SONET port, log in to AMI and enter the following parameters:

configuration port sonet timing <port> (network | internal)

The <port> variable indicates the SONET port that is to be modified and the (network|internal) variable designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source. If there are no SONET ports on the switch fabric, then this option is disabled. For more information about configuring the internal timing source, please refer to the section on Network Module Configuration in Appendix B of this manual.

After the clocking source is changed on a SONET port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

CHAPTER 5 DS-3 Configuration

If your Cabletron switch is equipped with one or more DS-3 network modules, some additional configuration may be necessary. There are several parameters which can be configured in order to work with other DS-3 equipment or to perform testing on the switch's DS-3 ports.

5.1 DS-3 Front Panel LEDs

There is an LED corresponding to both the transmit and receive lines of each DS-3 port. These LEDs provide some information about the state of the port depending on their color.

5.1.1 Transmit Indicators

The LED corresponding to the transmit line of the port has the following meaning depending on its color:

off Indicates that there are currently no cells being transmitted from the port.

green Indicates that a cell is being transmitted on the port.

5.1.2 Receive Indicators

The LED corresponding to the receive line of the port has the following meaning depending on the color:

off Indicates that a carrier has been detected on the line. A carrier is detected when there is a proper voltage signal on the line.

green Indicates that a cell is being received on the port.

red Indicates a loss of carrier.

5.2 Configuring DS-3 Mode

There are two modes of operation that the DS-3 ports support: PLCP (physical layer convergence protocol) and HCS (header check sequence). These modes control the way ATM cells are constructed from the DS-3 data stream.

To change the mode of operation of a DS-3 port, log in to the ATM Management Interface (AMI). (Please refer to Appendix A of this manual for information about logging into AMI). Enter the following parameters:

configuration port ds3 mode <port> (plcp | hcs)

The <port> variable indicates the DS-3 port that is to be modified and the (plcp|hcs) variable is the mode of operation to be used for the port. It is critical that both ends of the DS-3 link be configured with the same mode, regardless of which one is chosen. If there are no DS-3 ports on the switch fabric, then this option is disabled.

After the mode configuration is changed, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the new configuration will be put into effect every time the switch control software starts on that particular SCP.

5.3 Configuring DS-3 Empty Cells

To change the type of cells sent as empty cells (filler that is sent when a port is not sending data) on a DS-3 port, log in to AMI and enter the following parameters:

configuration port ds3 emptycells <port> (idle | unassigned)

The <port> variable indicates the DS-3 port that is to be modified and the (idle | unassigned) variable indicates the type of cells that the specified port sends as filler when the port is not sending data. The default setting is unassigned. If there are no DS-3 ports on the switch fabric, then this option is disabled.

After the type of empty cells is modified on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

5.4 Configuring DS-3 Line Length

To change the line length of a DS-3 port to correspond to the physical cable attached to that port, log in to AMI. Enter the following parameters:

configuration port ds3 length <port> (Lt225 | Gt225)

The <port> variable indicates the DS-3 port that is to be modified and the (Lt225 | Gt225) variable indicates the length of the cable attached to the port. Lt225 means the cable is less than 225 feet and Gt225 means the cable is greater than 225 feet. If there are no DS-3 ports on the switch fabric, then this option is disabled.

After the line length is modified on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

5.5 Configuring DS-3 Framing

There are two types of framing that the DS-3 ports support: clear channel and C-bit parity. These types of framing affect the significance of the overhead bits in the DS-3 frames.

To change the framing of a DS-3 port, log in to AMI and enter the following parameters:

configuration port ds3 framing <port> (cchannel | cbit)

The <port> variable indicates the DS-3 port that is to be modified and the (cchannel|cbit) variable is the type of framing for the port. The default setting is cbit (cbitparity). If there are no DS-3 ports on the switch fabric, then this option is disabled.

After the framing configuration is changed on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the new configuration will be put into effect every time the switch control software starts on that particular SCP.

5.6 Configuring DS-3 Payload Scrambling

Some equipment is more sensitive than the 9A000, SFCS-200BX or the SFCS-1000 when synchronizing with the network clock. This equipment requires a greater transition density than the switch fabric. To work with such equipment, it may be necessary to enable payload scrambling on the switch fabric. If a DS-3 port is configured for HCS cell delineation, scrambling should be enabled on both ends of the connection.

When payload scrambling is enabled, a scrambling function[1] is applied to the 48-byte payload of each cell transmitted. To operate with other equipment, scrambling must be enabled on the other end of the DS-3 connection as well so that the data is properly unscrambled. By default, scrambling is disabled.

To enable or disable scrambling on a DS-3 port, log in to AMI and enter the following parameters:

configuration port ds3 scrambling <port> (on | off)

Using the `on` variable means that cell payload scrambling is enabled on the specified port. Using the `off` variable means that cell payload scrambling is disabled on the specified port. Only the payload of the ATM cells is scrambled.

After scrambling is enabled or disabled on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

1. Cabletron switches use the recommended self-synchronizing scrambler, $x^{43} + 1$.

5.7 Configuring DS-3 Loopback

In order to facilitate testing of the DS-3 ports, there are four different loopback configurations available: cell, payload, diagnostic, and line. When a DS-3 port is in loopback mode, it no longer passes normal traffic.

To change the loopback state on a DS-3 port, log in to AMI and enter the following parameters:

configuration port ds3 loopback <port> (cell | payload | diag | line | none)

The <port> variable indicates the DS-3 port that is to be modified and the (cell | payload | diag | line | none) variable indicates the type of loopback to be used on the specified port. The default loopback setting is none, which means that no loopback will take place on that port. If there are no DS-3 ports on the switch fabric, then this option is disabled.

After the loopback state is modified on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

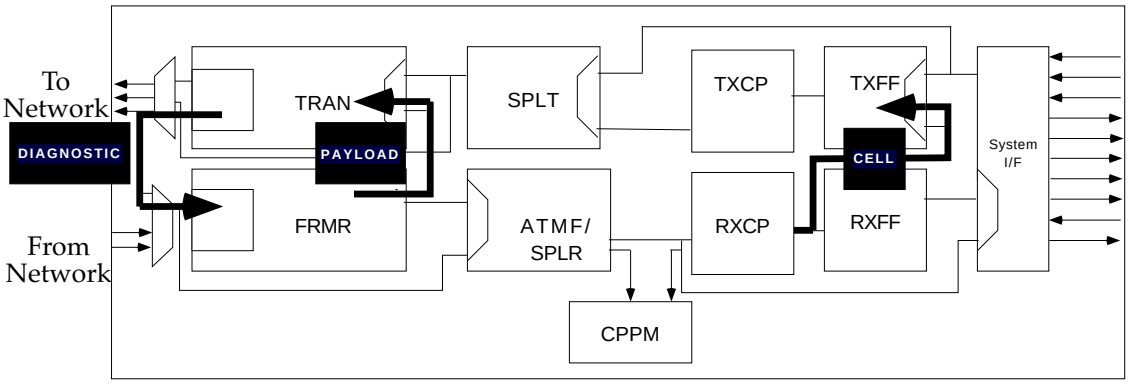

Figure 5.1 - DS-3 Single Port Loopback Diagram

5.7.1 Cell Loopback

When enabled, the DS-3 stream is received from the network, unframed into ATM cells, reframed, and then transmitted back to the network.

5.7.2 Payload Loopback

When enabled, the DS-3 stream is received from the network, has the DS-3 overhead bits re-inserted, and is retransmitted to the network.

5.7.3 Diagnostic Loopback

This connects the receiver to the transmitter. The DS-3 stream transmitted by the switch to a port is looped back to the switch. The DS-3 stream is still transmitted to the network, but the incoming DS-3 stream is ignored.

5.7.4 Line Loopback

Line loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted back out to the line. Cells generated by the switch to this port are not sent over the line.

5.8 Displaying DS-3 Error Counters

The user can display several DS-3 counters by logging in to AMI. Enter the following parameters to display the DS-3 network module statistics:

statistics ds3

```
ds3 Port 1C1 Counter            Value           Delta
------------------------------  --------------  --------------
ds3FramingLOSs                               0               0
ds3FramingLCVs                               0      3047081531
ds3FramingSumLCVs                   3533754755               0
ds3FramingFERRs                      122265891               0
ds3FramingOOFs                           59758               0
ds3FramingFERFs                              0               0
ds3FramingAISs                               0               0
ds3FramingPbitPERRs                 1713276195               0
ds3FramingCbitPERRs                 1583241699               0
ds3FramingFEBEs                      976095339               0
ds3PlcpFERRs                            956136               0
ds3PlcpLOFs                                  0               0
ds3PlcpBIP8s                         252419904               0
ds3PlcpFEBEs                         603317432               0
ds3PlcpYellows                               0               0
ds3AtmHCSs                           176407092               0
ds3AtmRxCells                       3709807680               0
ds3AtmTxCells                           203023               0

Press return for more, q to quit: q
```

All of the PLCP counters listed above and the Yellow counter have meaningful values only when the DS-3 network module is running in the PLCP mode. They are all meaningless when running in the HCS mode.

However, the HCS counter always has meaning, regardless of which mode is running.

5.9 DS-3 Error Counter Descriptions

ds3FramingLOSs Indicates the number of seconds in which Loss of Signal (LOS) errors were detected by the DS3 Receive Framer block.

ds3FramingLCVs Shows the number of Line Code Violations (LCV) that were detected by the DS3 Receive Framer block.

ds3FramingSumLCVs Designates the number of DS3 information blocks (85 bits) which contain one or more Line Code Violations (LCV).

ds3FramingFERRs Lists the number of DS3 framing error (FERR) events.

ds3FramingOOFs Indicates the number of DS3 Out Of Frame (OOF) error events.

ds3FramingFERFs Shows the number of seconds in which Far End Receive Failure (FERF) state has been detected by the DS3 Receive Framer block. The FERF signal alerts the upstream terminal that a failure has been detected along the downstream line.

ds3FramingAISs Displays the number of seconds in which Alarm Indication Signals (AIS) were detected by the DS3 Receive Framer block. AIS indicates that an upstream failure has been detected by the far end.

ds3FramingPbitPERRs Lists the number of P-bit parity error (PERR) events.

ds3FramingCbitCERRs Indicates the number of C-bit parity error (CERR) events.

ds3FramingFEBEs Designates the number of DS3 far end block error (FEBE) events.

ds3PlcpFERRs Displays the number of Physical Layer Convergence Protocol (PLCP) octet error events.

ds3PlcpLOFs Shows the number of seconds in which Loss Of Frame (LOF) errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. LOF is declared when an Out-Of-Frame state persists for more than 1ms. LOF is removed when an in-frame state persists for more than 12ms.

ds3PlcpBIP8s Displays the number of BIP-8 (Bit Interleaved Parity-8) error events. The BIP-8 is calculated over the Path Overhead field and the associated ATM cell of the previous frame. A BIP-N is a method of error monitoring. An N-bit code is generated by the transmitting equipment in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the previous VT SPE, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc.

ds3PlcpFEBEs Indicates the number of ATM Far End Block Error (FEBE) events.

ds3PlcpYellows Designates the number of seconds in which Yellow alarm errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. Yellow alarm is asserted when 10 consecutive yellow signal bits are set to logical 1. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

ds3AtmHCSs Displays the number of header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

ds3AtmRxCells Lists the number of ATM cells that were received, not including idle/unassigned cells.

ds3AtmTxCells Shows the number of non-null ATM cells that were transmitted, not including idle/unassigned cells.

5.10 Configuring DS-3 Timing

The DS-3 ports on an individual switch fabric can derive timing from one of two sources: an internal clock (internal) or the incoming DS-3 data stream (network). The default setting for ports on the switch is internal clocking. Cabletron Systems recommends that the following clocking conventions be used to configure the DS-3 clocking on an individual switch fabric:

- When connecting to a carrier-provided DS-3 service, the carrier's recommendation should be followed. In most cases, the carrier provides a clock and the switch should be configured for network timing.
- When connecting two Cabletron switches over a cable within the campus, both ends of the connection should be set to internal clocking (default).
- When connecting two Cabletron switches through other DS-3 equipment within the campus, it may be necessary to change the clocking to network (depending on the type of equipment). In this case, the recommendation of the campus network administration should be followed.

To change the clocking source for a DS-3 port, log in to AMI and enter the following parameters at the prompt to change the clocking source:

configuration port ds3 timing <port> (network | internal)

The <port> variable indicates the DS-3 port that is to be modified and the (network|internal) variable designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source. If there are no DS-3 ports on the switch fabric, then this option is disabled. For more information about how to configure the internal timing source, please refer to the section on Network Module Configuration in Appendix B of this manual.

After the clocking source is changed on a DS-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

CHAPTER 6 E-3 Configuration

E-3 network modules may require some additional configuration. There are several parameters that can be configured in order to work with other E-3 equipment or to perform testing on the E-3 ports.

6.1 E-3 Front Panel LEDs

There is an LED corresponding to both the transmit and receive lines of each E-3 port. These LEDs provide some information about the state of the port depending on their color.

6.1.1 Transmit Indicators

The LED corresponding to the transmit line of the port has the following meaning depending on its color:

off Indicates that there are currently no cells being transmitted from the port.

green Indicates that a cell is being transmitted on the port.

6.1.2 Receive Indicators

The LED corresponding to the receive line of the port has the following meaning depending on the color:

off Indicates that a carrier has been detected on the line. A carrier is detected when there is a proper voltage signal on the line.

green Indicates that a cell is being received on the port.

red Indicates a loss of carrier.

6.2 Configuring E-3 Mode

There are two modes of operation that the E-3 ports support: PLCP (also referred to as G.751) and HCS (also referred to as G.832). These modes control the way ATM cells are delineated within the E-3 data stream.

To change the mode of operation of an E-3 port, log in to the ATM Management Interface (AMI). (Please refer to Appendix A of this manual for information about logging into AMI). Enter the following parameters to enable PLCP or HCS framing:

configuration port e3 mode <port> (plcp | hcs)

The <port> variable indicates the E-3 port that is to be modified and the (plcp|hcs) variable is the mode of operation to be used for the port. It is critical that both ends of the E-3 link be configured with the same mode, regardless of which one is chosen. hcs is the default. If there are no E-3 ports on the switch fabric, then this option is disabled.

After the E-3 mode configuration is changed, the switch control processor (SCP) immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

6.3 Configuring E-3 Payload Scrambling

Some equipment is more sensitive than the switch when synchronizing with the network clock. This equipment requires a greater transition density than the switch. To work with such equipment, it may be necessary to enable payload scrambling on the SCPs.

When payload scrambling is enabled, a scrambling function[1] is applied to the 48-byte payload of each cell transmitted. To operate with other equipment, scrambling must be enabled on the other end of the E-3 connection as well so that the data is properly unscrambled. By default, scrambling is disabled.

1. Cabletron switches use the recommended self-synchronizing scrambler, $x^{43} + 1$.

To enable or disable scrambling on an E-3 port, log in to AMI and enter the following parameters:

configuration port e3 scrambling <port> (on | off)

Using the on variable means that cell payload scrambling is enabled on the port. Using the off variable means that cell payload scrambling is disabled on the port. Only the payload of the ATM cells is scrambled. If there are no E-3 ports on the switch fabric, then this option is disabled.

After scrambling is enabled or disabled on an E-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

6.4 Configuring E-3 Loopback

To facilitate testing of the E-3 ports, there are four different loopback configurations available: cell, payload, diagnostic, and line. When an E-3 port is in loopback mode, it no longer passes normal traffic.

To change the loopback state on an E-3 port, log in to AMI and enter the following parameters:

configuration port e3 loopback <port> (cell | payload | diag | line| none)

The <port> variable indicates the E-3 port that is to be modified and the (cell | payload | diag | line | none) variable indicates the type of loopback on the port. The default loopback setting is none which means that no loopback will take place on that port. If there are no E-3 ports on the switch fabric, then this option is disabled.

After the loopback mode is modified on a E-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

.

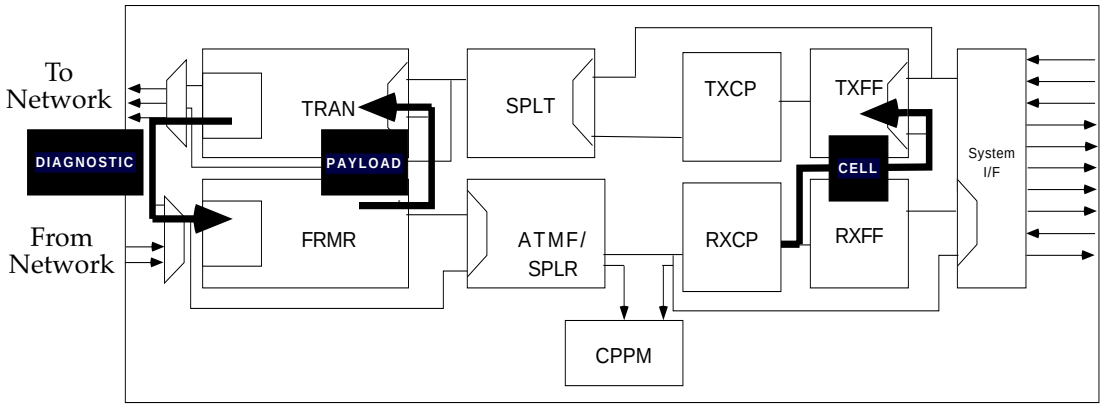

Figure 6.1 - E-3 Loopback

6.4.1 Cell Loopback

When enabled, the E-3 stream received from the network is unframed into ATM cells. The cells are then reframed and transmitted back to the network.

6.4.2 Payload Loopback

When enabled, the E-3 stream received from the network has the E-3 overhead bits re-inserted and is retransmitted to the network.

6.4.3 Diagnostic Loopback

This connects the receiver to the transmitter. The E-3 stream transmitted by the SCP to a port are looped back to the SCP. The E-3 stream is still transmitted to the network, but the incoming E-3 stream is ignored.

6.4.4 Line Loopback

Line loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted back out to the line. Cells generated by the SCP to this port are not sent over the line.

6.5 Configuring E-3 Empty Cells

To change the type of cells sent as empty cells (filler that is sent when a port is not sending data) on an E-3 port, enter the following parameters in AMI:

configuration port e3 emptycells <port> (idle | unassigned)

The <port> variable indicates the E-3 port that is to be modified and the (idle | unassigned) variable indicates the type of cells that the specified port sends as filler when the port is not sending data. The default setting is unassigned. If there are no E-3 ports on the switch, then this option is disabled.

After the type of empty cells is modified on a E-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

6.6 Displaying E-3 Error Counters

The user can display several E-3 counters by logging in to AMI. Enter the following parameters to display the E-3 network module statistics:

statistics e3

```
e3 Port 1D1 Counter                Value          Delta
-----------------------------  -------------  -------------
e3FramingLOSs                          85974              0
e3FramingLCVs                     3684415794              0
e3FramingFERRs                      85173622              0
e3FramingOOFs                          85974              0
e3FramingFERFs                             0              0
e3FramingAISs                              0              0
e3FramingBIP8s                     636877586              0
e3FramingFEBEs                       2465566              0
e3PlcpFERRs                                0              0
e3PlcpLOFs                            171950              0
e3PlcpBIP8s                                0              0
e3PlcpFEBEs                                0              0
e3PlcpYellows                              0              0
e3AtmHCSs                                  0              0
e3AtmRxCells                               0              0
e3AtmTxCells                          281929              0

Press return for more, q to quit: q
```

All of the PLCP counters listed above and the Yellow counter have meaningful values only when the E-3 network module is running with PLCP framing. They are all meaningless when running in the HCS mode.

However, the HCS counter always has meaning, regardless of which mode is running.

6.7 E-3 Error Counter Descriptions

e3FramingLOSs Specifies the number of seconds in which Loss Of Signal (LOS) errors were detected by the E3 Receive Framer block.

e3FramingLCVs Displays the number of Line Code Violations (LCV) that were detected by the E3 Receive Framer block.

e3FramingFERRs Shows the number of E3 framing error (FERR) events.

e3FramingOOFs Lists the number of E3 Out Of Frame (OOF) error events.

e3FramingFERFs Indicates the number of Far End Receive Failures for a port configured with HCS framing. Indicates the number of Remote Alarm Indications for a port configured with PLCP framing.

e3FramingAISs Specifies the number of seconds in which Alarm Indication Signals (AIS) were detected by the E3 Receive Framer block. AIS indicates that an upstream failure has been detected by the far end.

e3FramingFEBEs Displays the number of E3 far end block error (FEBE) events.

e3FramingBIP8s Shows the number of E3 G.832 BIP-8 errors. This counter is only valid for a port configured with HCS framing.

e3PlcpFERRs Lists the number of Physical Layer Convergence Protocol (PLCP) octet error events.

e3PlcpLOFs Indicates the number of seconds in which Loss Of Frame (LOF) errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. LOF is declared when an Out Of Frame state persists for more than 1ms. LOF is removed when an in-frame state persists for more than 12ms.

e3PlcpBIP8s Indicates the number of BIP-8 (Bit Interleaved Parity - 8) error events. The BIP-8 is calculated over the Path Overhead field and the associated ATM cell of the previous frame. A BIP-N is a method of error monitoring. An N-bit code is generated by the transmitting equipment in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the previous VT SPE, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc.

e3PlcpFEBEs Displays the number of ATM Far End Block Error (FEBE) events.

e3PlcpYellows Shows the number of seconds in which Yellow alarm errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. Yellow alarm is asserted when 10 consecutive yellow signal bits are set to logical 1. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

e3AtmHCSs Lists the number of header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

e3AtmRxCells Indicates the number of ATM cells that were received, not including idle/unassigned cells.

e3AtmTxCells Displays the number of non-null ATM cells that were transmitted, not including idle/unassigned cells.

6.8 Configuring E-3 Timing

The E-3 ports on an individual switch fabric can derive timing from one of two sources: an internal clock (internal) or the incoming E-3 data stream (network). The default is internal. Cabletron recommends that the following conventions be used to configure the E-3 clocking on an individual switch fabric:

- When connecting to a carrier-provided E-3 service, the carrier's recommendation should be followed. In most cases, the carrier provides a clock and the SCP should be configured for network timing.
- When connecting two Cabletron switch fabrics over a cable within the campus, both ends of the connection should be set to internal clocking (default).
- When connecting two Cabletron switch fabrics through other E-3 equipment within the campus, it may be necessary to change the clocking to network (depending on the type of equipment). In this case, the recommendation of the campus network administration should be followed.

To change the clocking source for a E-3 port, log in to AMI and enter the following parameters:

configuration port e3 timing <port> (network | internal)

The <port> variable indicates the E-3 port that is to be modified and the (network|internal) variable designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source. If there are no E-3 ports on the switch fabric, then this option is disabled. For more information about how to configure the internal timing source, please refer to the section on Network Module Configuration in Appendix B of this manual.

After the clocking source is changed on an E-3 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

CHAPTER 7 TP25 Configuration

If your switch fabric is equipped with one or more TP25 (25 Mbps) network modules, there is some additional configuration that may be necessary. There are several parameters which can be configured in order to work with other TP25 equipment or to perform testing on the TP25 ports.

7.1 TP25 Front Panel LEDs

There is an LED corresponding to both the transmit and receive lines of each TP25 port. These LEDs provide some information about the state of the port depending on their color.

7.1.1 Transmit Indicators

The LED corresponding to the transmit line of the port has the following meaning depending on its color:

off Indicates that there are currently no cells being transmitted from the port.

green Indicates that a cell is being transmitted on the port. The intensity of the green light increases as the traffic on the port increases.

7.1.2 Receive Indicators

The LED corresponding to the receive line of the port has the following meaning depending on the color:

off Indicates that a carrier has been detected on the line. A carrier is detected when there is a proper voltage signal on the line.

green Indicates that a cell is being received on the port. The intensity of the green light increases as the traffic on the port increases.

red Indicates a loss of carrier.

7.2 Configuring TP25 Loopback

To facilitate testing of the TP25 ports, a line loopback configuration is available. When a TP25 port is in loopback mode, it no longer passes normal traffic.

To change the loopback state on a TP25 port, log in to AMI. (Please refer to Appendix A of this manual for information about logging into AMI). Enter the following parameters:

configuration port tp25 loopback <port> (line |none)

The <port> variable indicates the TP25 port that is to be modified and the (line|none) variable indicates the type of loopback on the port. The default loopback setting is none which means that no loopback will take place on that port. If there are no TP25 ports on the switch, then this option is disabled.

After the loopback mode is modified on a TP25 port, the SCP immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file so that the change will be put into effect every time the switch control software starts on that particular SCP.

7.2.1 Line Loopback

Line loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted back out to the line. Cells generated by the switch to this port are not sent over the line.

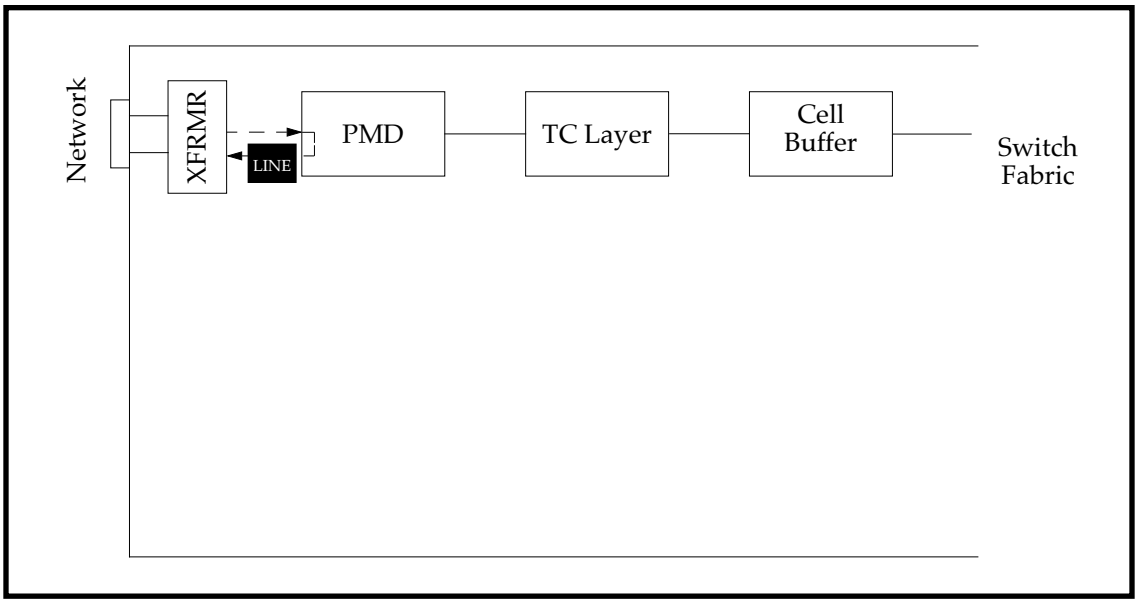

Figure 7.1 - TP25 Single Port Loopback Diagram

7.3 Displaying TP25 Error Counters

The user can display several TP25 counters by logging in to AMI. Enter the following parameters to display the TP25 network module statistics:

statistics tp25

```
tp25 Port 1A1 Counter               Value           Delta
----------------------------  --------------  --------------
tp25ErrorSymbol                   40452300               0
tp25AtmHCSs                             8                0
tp25AtmRxCells                      13722                0
tp25AtmTxCells                          0                0

Press return for more, q to quit: q
```

7.4 TP25 Error Counter Descriptions

tp25ErrorSymbols Lists the number of undefined symbols received.

tp25AtmHCSs Indicates the number of header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

tp25AtmRxCells Displays the number of ATM cells that were received.

tp25AtmTxCells Shows the number of ATM cells that were transmitted.

CHAPTER 8 TAXI Configuration

The TAXI network modules for the Cabletron ATM switches use the 4B/5B TAXI format and are shipped preconfigured. No additional configuration is needed for the TAXI network modules to interoperate with another Cabletron Systems ATM switch or with any other 4B/5B TAXI interface.

8.1 TAXI Front Panel LEDs

There is an LED corresponding to both the transmit and receive lines of each TAXI port. These LEDs provide some information about the state of the port depending on their color.

8.1.1 Transmit Indicators

The LED corresponding to the transmit line of the port has the following meaning depending on its color:

off Indicates that there are currently no cells being transmitted from the port.

green Indicates that a cell is being transmitted on the port. The intensity of the green light increases as the traffic on the port increases.

8.1.2 Receive Indicators

The LED corresponding to the receive line of the port has the following meaning depending on its color:

off Indicates that a carrier has been detected on the line. A carrier is detected when there is a proper optical signal on the line.

green Indicates that a cell is being received on the port. The intensity of the green light increases as the traffic on the port increases.

red Indicates a loss of carrier.

8.2 Configuring TAXI Loopback

To facilitate testing of the TAXI ports, a diagnostic loopback configuration is available. When a TAXI port is in loopback mode, it will no longer pass normal traffic.

To change the loopback state on a TAXI port, log in to the ATM Management Interface (AMI) and open a session. (Please refer to Appendix A of this manual for information about logging into AMI.) Enter the following parameters:

configuration port taxi loopback <port> (diag | none)

The <port> variable indicates the TAXI port that is to be modified and the (diag|none) variable indicates the type of loopback to be used on the specified port. The default loopback setting is none, which means that no loopback will take place on that port. If there are no TAXI ports on the switch fabric, then this option is disabled.

After the loopback state is modified on a TAXI port, the switch control processor (SCP) immediately puts that change into effect on that switch fabric. At the same time, that information is entered into the configuration database file on the SCP so that the change will be put into effect every time the switch control software starts on that particular SCP.

APPENDIX A AMI Overview

The switch software provides switch and connection management, IP connectivity, and SNMP network management. The Switch Control Software (SCS) is the "brains" of the switch. The SCS controls the switch board(s) and handles connection set-up and tear-down duties. The SCS can also communicate with other FORE Systems switches using the SPANS NNI protocol to learn network topology and establish connections across multiple switches. In addition, there is an SNMP agent built into the SCS to allow SNMP management and control.

The user interface to the SCS is called the ATM Management Interface (AMI). AMI can be run on any Cabletron switch running switch software version 3.0.1 or later. This chapter contains a description of how to log in to AMI, how to open or close an AMI session, and how to perform other AMI root level commands. AMI allows the user to configure and to make statistical queries of various hardware and software aspects of Cabletron switches and network modules by providing a hierarchical menu system similar to a UNIX file system. A single root menu provides a number of commands. Some of those commands, in turn, call submenus which provide a number of subcommands. At any given time, the user works within a particular submenu which is indicated by the prompt. The user can traverse a submenu one level at a time, or can traverse a number of levels simultaneously if the entire command string is known. For example, to show the current configuration of the network modules, enter the following at the prompt:

```
localhost::> configuration module show
```

rather than entering one command line at a time as follows:

```
localhost::> configuration

localhost::configuration> module

localhost::configuration module> show
```

Additionally, the user only needs to enter the minimum number of letters in a command which would make the command unique to that level. For example, the user could enter `co m s` instead of `configuration module show`. However, the minimum number of letters entered must also distinguish the command from global commands, such as `top` or `up`. For example, you would have to enter `topo` to distinguish `topology` from the global command `top` or `upc` to distinguish `upc` from the global command `up`.

AMI is described throughout this chapter using the following conventions:

- All AMI output, including user prompts, is shown in `courier` font.
- All user input; e.g., sub-commands, is shown in `avant garde` font.
- Each submenu is described in a separate section.
- Commands that are submenus are immediately followed by a ">" symbol. The ">" should not be entered as part of the command.
- Required parameter values are indicated inside angle brackets "< >".
- Optional parameter values are indicated inside square brackets "[ ]". The "[ ]" should not be entered as part of the command.
- Parameter values that require a choice are separated by vertical bars and are enclosed in parentheses "(|)" Neither the vertical bar nor the parentheses should be entered as part of the command.
- Optional parameter names are indicated with dashes "-".
- All port numbers are in BNP (board, network module, port) notation.

A.1 Initial Login from Serial Port or Telnet

The user can log in to the switch either through the serial port or through the Ethernet port using telnet.

A.1.1 Login from Serial Port

When connecting to the switch via the serial port, output similar to the following will be displayed on your console:

```
ForeThought_4.0.0 (1.15) (asx200bx) (fishtank)
```

Above, ForeThought_4.0.0 (1.15) indicates the version of software, (asx200bx) indicates what type of switch this is, and (fishtank) indicates the name that has been assigned to this SCP. If (ATM SWITCH) is displayed for the switch name, this means that no host name has been assigned yet.

At the login prompt, if a password has been assigned to the switch, the user should type asx <ENTER> and then type the password. The switch will not echo the keystrokes of the password for security reasons.

```
login: asx
Password:
```

If no password has been assigned, enter asx <ENTER> at the serial port. In either case, the following is displayed and a session is opened on the SCP:

```
ATM Management Interface v1.2
Copyright (c) 1994, 1995 FORE Systems, Inc.
All Rights Reserved

General commands:
  '?' to get list of commands at the current level
  'up' to go up one menu level
  'top' to go to the root menu
  'exit' to leave AMI

Opening a session for “127.0.0.1”, please wait...

Connected to “127.0.0.1” (asx200bx).

localhost::>
```

A.1.2 Login from Telnet

To telnet to the SCP, enter the following parameters at the > prompt on the host:

```
> telnet <name>
```

name Enter either the name or the IP address of the SCP.

For example, to telnet to an SCP called fishtank, enter the following:

```
> telnet fishtank
```

When the telnet connection is established, something similar to the following will be displayed:

```
Trying 204.95.89.231 ...
Connected to fishtank.
Escape character is '^]'.

ForeThought_4.0.0 (1.15) (asx200bx) (fishtank)
```

Above, ForeThought_4.0.0 (1.15) indicates the version of software, (asx200bx) indicates what type of switch this is, and (fishtank) indicates the name that has been assigned to this SCP. If (ATM SWITCH) is displayed for the switch name, this means that no host name has been assigned yet.

Only one user may open an AMI session on an SCP at a time. The user will be prompted to log in. The user should enter ami at the prompt.

If a password has been assigned, then the user will be prompted for that password. The keystrokes of the password will not be echoed for security reasons. If no password has been assigned (e.g., the very first time you log in), then the user will not be prompted for a password. For more information about how to assign or to change the password, please refer to the section entitled "Setting or Changing the Password" under "Operation Commands" in Appendix C of this manual.

```
login: ami
Password:
```

The following is displayed and a session is opened on the SCP:

```
ATM Management Interface v1.2
Copyright (c) 1994, 1995 FORE Systems, Inc.
All Rights Reserved


General commands:
  '?' to get list of commands at the current level
  'up' to go up one menu level
  'top' to go to the root menu
  'exit' to leave AMI


Opening a session for “127.0.0.1”, please wait...

Connected to “127.0.0.1” (asx200bx).

localhost::>
```

If another user already has an AMI session open on that SCP, then you will not be permitted to log in and will receive the following message:

```
Another ami is running on this switch. . Exiting...
Connection closed by foreign host.
```

A.2 AMI Commands Not Available When Running Remotely

Some AMI commands are not available when you telnet or log in to a switch remotely. For example, if you are logged in locally to a switch called fishtank (you will see `localhost::>` as your prompt) and you open a session to a switch called shark (you will see `shark::>` as your prompt), there are some AMI commands that will not work on shark.

The following is a list of the commands that are not available when running a remote AMI session, and the applicable switch platforms:

For the 9A000, SFCS-200BX, SFCS-200WG, and SFCS-1000 platforms:

- configuration lane lecs get
- operation cdb backup
- operation cdb restore
- operation password
- operation reboot
- operation upgrade

For the 9A000, SFCS-200WG, SFCS-200BX, and SFCS-1000 platforms:

- configuration system syslog
- configuration system timeout
- operation cdb init
- operation flash
- operation panic
- operation version

A.3 AMI Root Menu for an Open Session

This menu is the root submenu for an AMI session. When AMI is first entered from the serial port or telnet, the localhost session is the only open session. The following information is displayed:

```
ATM Management Interface v1.2
Copyright (c) 1994, 1995 FORE Systems, Inc.
All Rights Reserved

General commands:
  '?' to get list of commands at the current level
  'up' to go up one menu level
  'top' to go to the root menu
  'exit' to leave AMI

Opening a session for “localhost”, please wait...

Connected to “localhost” (asx200bx).

localhost::>
```

By typing a “?” at any prompt, a list of available commands at the current level is displayed. By typing a “?” at this root level prompt, the following command list is shown:

```
localhost::> ?

about            close            configuration>   exit
help             history          open             operation>
ping             redo             rows             statistics>
top              up
```

Each of these root level commands is described in the following subsections.

A.3.1 About Command

By entering the `about` command at the root level prompt, the user can display information regarding AMI and how to begin an AMI session on a host or on a switch.

```
localhost::> about

ATM Management Interface v4.0
Copyright (c) 1994, 1995 FORE Systems, Inc.
All Rights Reserved

AMI uses SNMP to manage FORE Systems' ATM switches. AMI is platform
independent and runs on hosts and FORE ATM switches. When AMI is
executed on a host, you must first use the OPEN command to
specify the switch to manage.  If AMI is started on the switch,
it immediately opens a connection to the local switch.
```

A.3.2 Close Command

Any number of sessions may be opened to remote SCPs from your local SCP. However, only one AMI session may be open at a time on any given SCP on a 9A000, SFCS-200WG, SFCS-200BX, or an SFCS-1000.

By typing `close` at the prompt, the user can end the current AMI session.

If an individual session is closed, the user is sent back to the last session that is still open. For example, if you opened a session on switch1 and on switch2 from your local SCP, and you wanted to close the session on switch2, you would be sent back to the last open session which is on switch1 as follows:

```
switch2::> close

switch1::>
```

If you decided to close the session on switch1, you would be sent back to the last open session which is on your local SCP as follows:

```
switch1::> close

localhost::>
```

If all sessions are closed, the user is sent back to the root prompt as follows:

```
localhost::> close

>
```

At this point, the user can open another session or exit the switch.

A.3.3 Configuration Commands

By entering configuration at the root level, the user can access several sub-commands that allow the user to configure specific parts of the hardware or specific properties of the software. The configuration commands are described in detail in Appendix B of this manual.

A.3.4 Exit Command

The exit command lets the user log out of the main AMI system. When entered, this command ends all open sessions on the SCP. Enter the following:

```
localhost::> exit
Connection closed by foreign host.
```

A.3.5 Help Command

By typing help at any submenu level, a list of available commands at the current level, and a short description of each command, is shown. By typing help at the root level, the following commands and descriptions are shown:

```
localhost::> help
General commands:
  '?' to get list of commands at the current level
  'up' to go up one menu level
  'top' to go to the root menu
  'exit' to leave AMI

  about                  - Display program information
  close                  - Close this connection
  configuration>         - System configuration submenu
  exit                   - Exit AMI
  help                   - Display help for each command
  history                - Display command history
  open                   - Open a connection
  operation>             - Switch operation submenu
  ping                   - Ping a host or switch
  redo                   - Repeat a history command
  rows                   - Get/set number of rows
  statistics>            - Switch statistics submenu
  top                    - Go to the root menu
  up                     - Go up 1 menu level
```

A.3.6 History Command

By typing history at any prompt, the user can list up to the last 20 previously typed commands for that particular session as follows:

```
localhost::> history
1  open fishtank
2  stat
3  ?
4  module
5  oc3
6  port
7  spans
8  tcp
9  udp
10  vcc
11  vpc
12  up
13  help
14  history
```

A.3.7 Open Command

The `open` command lets the user begin a session on a remote switch. At the prompt, enter the following parameters:

```
localhost::> open <host> [<community>]
```

host Indicates the host name of the remote switch on which the user wants to open a session.

community Enter the SNMP community string that indicates level of access that the user has on the switch. The default is `public`, which allows read-only access.

NOTE: Although the default SNMP community string is `public`, you must use the `private` SNMP community string if you wish to make any changes on the remote SCP (e.g., if you want to create an SPVC to that SCP).

For example, to log in to a remote switch named fishtank using the `public` community string, enter the following parameters:

```
localhost:> open fishtank public

Opening a session for “fishtank”, please wait...

Connected to “fishtank” (asx200wg).

fishtank::>
```

If another user already has an AMI session open on that SCP, then you will not be permitted to log in to that SCP. You will receive the following message:

```
Another ami is running on this switch. . Exiting...
Connection closed by foreign host.
```

A.3.8 Operation Commands

By entering `operation` at the root level prompt, the user can access several subcommands that allow the user to manage various parts of the switch. The `operation` commands are described in detail in Appendix C of this manual.

A.3.9 Ping Command

The `ping` command lets the user send a ping to another switch or a host to see if it is "alive," or reachable, by sending it an ICMP echo request and waiting for a response. The user can access this command by entering `ping` at the root level. Enter the following parameters:

```
localhost::> ping <IP-address>
```

IP-address Indicates the IP address of the host or switch to which the ping will be sent.

NOTE: The ping is always sent from the first switch or host on which AMI was originally started. For example, you are logged into switch A. From there, you open a session to switch B. If you enter the `ping` command while in your session on switch B, the ping is sent from switch A, NOT from switch B.

A.3.10 Redo Command

The redo command can be used in conjunction with the history command. It lets the user repeat a command that was given in the same open session. The user can access this command by entering redo at any level. To repeat the last command that was performed, enter redo with no additional parameters as follows:

```
localhost::> redo
```

To repeat a command given within the last 20 commands in the same open session, enter the following parameters:

```
localhost::> redo <command-number>
```

command-number This is the command and the number associated with that command which was previously performed by the switch during this same session. The user should enter the history command to list the previous commands and their associated numbers as shown in the following example.

Type history at the prompt to list the last 20 previously typed commands for that particular session as follows:

```
localhost::> history

1  open localhost
2  stat
3  ?
4  module
5  oc3
6  port
7  spans
8  tcp
9  udp
10  vcc
11  help
12  history
```

Then, to repeat a previously given command, type redo and the command number at the prompt. For example, to repeat command number 8, which is listing statistics for tcp, enter the following:

```
localhost::> redo 8

tcp Counter                     Value
------------------------------ --------------------
tcpActiveOpens                                    1
tcpPassiveOpens                                  49
tcpAttemptFails                                   0
tcpEstabResets                                    1
tcpCurrEstab                                      1
tcpInSegs                                     14060
tcpOutSegs                                     9967
tcpRetransSegs                                    0
```

A.3.11 Rows Command

The rows command allows users to set the number of rows that their terminal displays. Users can access this command by entering rows at the root level as follows:

```
localhost::> rows [<rows>]

Terminal Rows = 24
```

rows Indicates the number of terminal rows to be used.

A.3.12 Statistics Commands

By entering statistics at the root level, the user can access several commands that display operational performance and error information for the various hardware and software features of the switch and the network modules. The statistics commands are described in Appendix D of this manual.

A.3.13 Top Command

By entering top at any level, the user is sent to the root level of AMI. For example, if you are at the operation cdb level and you want to go directly to the root level, simply enter top at the prompt as follows:

```
localhost::operation cdb> top

localhost::>
```

A.3.14 Up Command

Entering up allows the user to go up one menu level. For example, if you are at the configuration port e3 level and you want to go one level above that to configuration port, simply enter up at the prompt as shown here.

```
localhost::configuration port e3> up

localhost::configuration port>
```

APPENDIX B AMI Configuration Commands

This chapter contains a detailed description of the AMI configuration commands. The main configuration menu can be found at the root level. There are several commands available under configuration. Commands that are submenus are immediately followed by a ">" symbol. Typing configuration ? at the prompt at the root level displays the configuration commands as follows:

```
localhost::> configuration ?

  alarms>          atmarp>          board>           ip>
  lane>            module>          nsap>            port>
  rs232>           snmp>            spans>           spvc>
  switch>          system>          topology>        uni30>
  upc>             vcc>             vpc>
```

Each of these commands has a submenu of commands which are described in the following subsections.

B.1 Alarms Configuration Commands

This submenu allows the user to configure alarms. The user can display the list of available subcommands by typing alarms ? at the configuration level.

```
localhost::configuration alarms> ?
show             enable           disable          reset
```

B.1.1 Displaying Alarm Conditions

This command lets the user display the status of all alarms. The linkFailed and spansFailed alarms are available on all Cabletron switches. The powerSupplyInputFailed and tempSensorOverTemp alarms are available on all switches. The powerSupplyOutputFailed alarm is only available on a 9A000, SFCS-200BX and an SFCS-1000. The fanBankFailed alarm is only available on an SFCS-1000. Enter the following parameters:

```
localhost::configuration alarms> show

AlarmType                     AlarmStatus    MinorAlarm   MajorAlarm
powerSupplyInputFailed        inactive       disabled     enabled
powerSupplyOutputFailed       inactive       disabled     enabled
fanBankFailed                 inactive       disabled     enabled
tempSensorOverTemp            inactive       disabled     enabled
linkFailed                    inactive       enabled      disabled
spansFailed                   inactive       enabled      disabled

Major alarm relay status: off
Minor alarm relay status: off
```

The fields in this display have the following meanings:

AlarmType Displays the name of the alarm.

AlarmStatus Shows whether the state of the alarm is active (alarming) or inactive (not alarming). An alarm is active if the underlying condition is detected. For power supplies, the input failed alarm condition is active if the input voltage is not within the nominal range for the supply. This does not necessarily mean that an output failure will result. A power supply output failure condition is active if any power supply is failing or if it is physically removed.

MinorAlarm Disabled means that this alarm type will not cause a minor alarm. Enabled means that this alarm type will cause a minor alarm.

MajorAlarm Disabled means that this alarm type will not cause a major alarm. Enabled means that this alarm type will cause a major alarm.

Major alarm relay status — Off means that no major alarms are currently active. On means that one or more major alarms are currently active. Look at the AlarmStatus field to see which condition is in a state of alarm.

Minor alarm relay status — Off means that no minor alarms are currently active. On means that one or more minor alarms are currently active. Look at the AlarmStatus field to see which condition is in a state of alarm.

B.1.2 Enabling an Alarm

This command lets the user enable an alarm. Enter the following parameters:

```
localhost::configuration alarms> enable (major | minor) <alarm type>
```

major | minor — Designates whether the alarm type causes a major alarm or a minor alarm when that condition occurs.

alarm type — Indicates the kind of alarm condition. Valid parameters are displayed in the AlarmType field when the command string configuration alarms show is entered at the prompt.

For example, to enable an overtemperature condition that is detected by the overtemperature sensor as a major alarm, enter the following parameters:

```
localhost::configuration alarms> enable major tempSensorOverTemp
```

To verify that the change has taken effect, you can display the alarms:

```
localhost::configuration alarms> show

AlarmType                    AlarmStatus    MinorAlarm   MajorAlarm
powerSupplyInputFailed       inactive       disabled     enabled
powerSupplyOutputFailed      inactive       disabled     enabled
fanBankFailed                inactive       disabled     enabled
tempSensorOverTemp           inactive       disabled     enabled
linkFailed                   inactive       enabled      disabled
spansFailed                  inactive       enabled      disabled

Major alarm relay status: off
Minor alarm relay status: off
```

B.1.3 Disabling an Alarm

This command lets the user disable an alarm. Enter the following parameters:

```
localhost::configuration alarms> disable (major | minor) <alarm type>
```

major | minor Designates whether the alarm type causes a major alarm or a minor alarm when that condition occurs.

alarm type Indicates the alarm condition. Valid parameters are displayed in the AlarmType field when the command string configuration alarms show is entered at the prompt.

For example, to disable a link failure as a minor alarm, enter the following parameters:

```
localhost::configuration alarms> disable minor linkFailed
```

To verify that the change has taken effect, you can display the alarms:

```
localhost::configuration alarms> show

AlarmType                    AlarmStatus    MinorAlarm   MajorAlarm
powerSupplyInputFailed       inactive       disabled     enabled
powerSupplyOutputFailed      inactive       disabled     enabled
fanBankFailed                inactive       disabled     enabled
tempSensorOverTemp           inactive       disabled     enabled
linkFailed                   inactive       disabled     enabled
spansFailed                  inactive       enabled      disabled

Major alarm relay status: off
Minor alarm relay status: off
```

B.1.4 Resetting an Alarm

This command lets the user reset either the linkFailed alarm, the spansFailed alarm, or both alarms. Enter the following parameters:

```
localhost::configuration alarms> reset (<alarm type> | all)
```

alarm type Indicates which alarm to reset. Enter either link or spans.

all Indicates that both the linkFailed and the spansFailed alarms will be reset.

For example, to reset the linkFailed alarm, enter the following parameters:

```
localhost::configuration alarms> reset link

Alarm linkFailed reset.
```

To verify that the change has taken effect, you can display the alarms:

```
localhost::configuration alarms> show

AlarmType                    AlarmStatus    MinorAlarm   MajorAlarm
linkFailed                   inactive       enabled      disabled
spansFailed                  active         enabled      disabled

Major alarm relay status: off
Minor alarm relay status: on
```

B.2 ATM ARP Configuration Commands

These commands allow the user to manage the ATM ARP (address resolution protocol) features. The user can display the list of available subcommands by typing `atmarp ?` at the `configuration` level.

```
localhost::configuration> atmarp ?
arpserver>       delete          flush           getnsap
mapnsap          newclassicalip  newforeip       show
```

B.2.1 ARP Server Configuration Commands

These commands let the user configure the RFC-1577 ATM ARP server. The user can get to this level by entering `arpserver` at the configuration level. By entering `arpserver ?` at this level, the list of available subcommands for `arpserver` is displayed.

```
localhost::configuration atmarp> arpserver ?
  show              set
```

B.2.1.1 Displaying the ARP Server

This command shows the RFC-1577 ATM ARP server for the IP network. Enter the following parameters:

```
localhost::configuration atmarp arpserver> show <interface>

qaa0 ARP server address: 0x47.0005.80.ffe100.0000.f124.00de.0020481900de.00

qaa0 is enabled for ARP server service
```

interface Enter the name of one of the classical IP interfaces, such as `qaa0`, `qaa1`, `qaa2`, or `qaa3`. The default is `qaa0`, so that if no interface is entered, the switch shows the ARP server address for `qaa0`.

B.2.1.2 Setting the ARP Server Address

This command allows the user to set the address of the RFC-1577 ATM ARP server. Enter the following parameters:

```
localhost::configuration atmarp arpserver> set <NSAPaddress> [<interface>]
```

NSAPaddress Indicates the ATM network layer address for the RFC-1577 ATM ARP server.

interface Enter the name of one of the Classical IP interfaces, such as `qaa0`, `qaa1`, `qaa2`, or `qaa3`. The default is `qaa0`.

The switch itself can be used as an ARP server. To do this, set the ARP server address to be the NSAP address of the switch's control port.

B.2.2 Deleting an ARP Entry

This command allows the user to remove an ARP entry from the ATM ARP cache. Enter the following parameters:

```
localhost::configuration atmarp> delete <host>
```

host Indicates the host name or the IP address of the end-station for which the outgoing ARP entry is to be deleted.

NOTE: If you have ILMI enabled on your switch, ILMI creates an ATM ARP cache entry for each address that it registers. These entries cannot be deleted using this command.

B.2.3 Flushing the ATM ARP Cache

This command enables the user to delete the contents of the ATM ARP cache. Only dynamic ARP cache entries are removed. The switch asks the user to verify that flushing the ARP cache is the desired action. Enter the following parameters:

```
localhost::configuration atmarp> flush

Flush the ATM ARP cache [n]? n
```

To cancel the command, type n and press <ENTER>, or simply type <ENTER>.

To flush the ARP cache, type y and press <ENTER>.

B.2.4 Getting the NSAP Address for a Classical IP Interface

This command displays the NSAP address of a Classical IP interface. Enter the following parameters:

```
localhost::configuration atmarp> getnsap [<interface>]

qaa0 NSAP address: 47000580ffe1000000f12400de0020481900de00
```

interface Indicates the name of the Classical IP interface to be displayed. If no interface is specified, the NSAP address of qaa0 is displayed.

B.2.5 Creating an IP to NSAP Address Mapping

This command allows the user to create an ATM ARP cache entry mapping a particular IP address to its corresponding NSAP address. Enter the following parameters:

```
localhost::configuration atmarp> mapnsap <host> <NSAPaddress> [<interface>]
```

host Designates the IP address.

NSAPaddress Designates the NSAP address.

interface Designates the Classical IP interface that should be used to open connections to this NSAP address.

B.2.6 Creating a Classical IP PVC

This command allows the user to create a new Classical IP PVC ARP entry. All data is sent LLC/SNAP encapsulated. Enter the following parameters:

```
localhost::configuration atmarp> newclassicalip <host> <vpi> <vci> [<interface>]
```

host Indicates the host IP address of the remote IP endstation.

vpi Indicates the virtual path number of the Classical IP PVC.

vci Indicates the virtual channel number of the Classical IP PVC.

interface Indicates the name of the Classical IP interface to be used for this connection. The default is `qaa0`.

B.2.7 Creating a FORE IP PVC ARP Entry

This command enables the user to create a FORE IP PVC ARP entry. Data on this PVC is encapsulated using null encapsulation (also known as VC-based multiplexing) as specified in RFC-1483. Enter the following parameters:

```
localhost::configuration atmarp> newforeip <host> <vpi> <vci> (4|5) [<interface>]
```

host Indicates the IP address of the remote host.

vpi Indicates the virtual path number. This must be 0.

vci Indicates the virtual channel number.

4|5 Designates the connection's ATM Adaptation Layer (AAL) type.

interface Indicates the FORE IP interface to be used. The default is `asx0`.

B.2.8 Displaying the ATM ARP Entries

This command displays the current ATM ARP cache. Enter the following parameters:

```
localhost::configuration atmarp> show

localhost::configuration atmarp>
```

When the prompt is returned with no information displayed, as shown above, then the ATM ARP cache is empty.

The following is an example of an ATM ARP cache.

```
localhost::configuration atmarp> show
IPaddress        If    VPI   VCI   AAL    Type              Direction
198.29.22.9      asx0  0     63    aal5   foreIpSVC         pending
198.29.22.15     asx0  0     231   aal5   foreIpSVC         pending
198.29.22.37     asx0  0     65    aal34  foreIpSVC         pending
IPaddress        If    NSAP Address
198.29.17.3      qaa0  0x47.0005.80.ffe100.0000.f21b.0138.002048102754.00
198.29.17.10     qaa0  0x47.0005.80.ffe100.0000.f21b.0137.002048100be6.00
198.29.17.15     qaa0  0x47.0005.80.ffe100.0000.f21b.0137.00204810048d.00
198.29.17.52     qaa0  0x47.0005.80.ffe100.0000.f21b.0138.0020481b0138.00
```

The fields in this display have the following meanings:

IPaddress Indicates the IP address for this connection.

If Shows the name of the IP interface for this connection.

VPI Displays the virtual path number.

VCI Displays the virtual channel number.

AAL Displays the AAL type of the given connection.

Type Lists the kind of connection. Can be foreIpPVC, foreIpSVC, classicalIpPVC, or classicalIpSVC.

Direction Pending means that a connection has not (yet) been established. Incomplete means that the IP-to-ATM address mapping is not yet known for the given IP address.

NSAP Address Lists the NSAP address for this connection.

B.3 Switch Board Configuration Commands

This submenu allows the user to configure default values for the switch board. The user can display the list of available subcommands by typing `board ?` at the `configuration` level.

```
localhost::configuration> board ?
 clockscale       show             topology
```

B.3.1 Configuring the Clock Scaling Factor on a Switch Board

This is an advanced command that allows the user to set the clock scaling factor for traffic policing on a switch board. The switch board represents units of time (i.e., burst tolerances) in internal units of clock ticks. Because of the 40 MHz clock used on the switch fabric, the maximum unit of time that can be used by the traffic policing hardware is 0.838 seconds. However, the burst tolerance of some VBR connections is larger than 0.838 seconds, so they can not be correctly policed by the switch. To accommodate these cases, it is possible to scale the clock to represent larger amounts of time with the same number of ticks. The user must enter the following parameters:

NOTE: This command is only useful when performing traffic policing on VBR traffic with an extremely large burst tolerance that is destined for a very low speed link. If these conditions are not represented in your network, it is recommended that you leave this value at the default setting.

```
localhost::configuration board> clockscale <board> <factor>
```

board Indicates the number of the board being managed.

factor Indicates the factor by which you want to scale the clock. The scaling factor multiplied by 0.838 seconds should be greater than the burst tolerance of the connection. The default is 1.

B.3.2 Displaying the Board Configuration

This command shows the current configuration of the switch board. Enter the following parameters:

```
localhost::configuration board> show

Board Version Model       S/N    NMs VPIerrors VCIerrors ClockScale MulticastMode
1     1.0     asx200bx    4465   4        2665      1469          1 extended
```

The fields in this display are defined as follows:

Board This is the number of the board (switch fabric) being managed.

Version Designates the hardware version number.

Model Displays what type of board this is.

S/N Indicates the serial number.

NMs Shows the number of installed network modules.

VPIerrors Lists the VPI lookup error cells arriving at the switch on an unknown VPI.

VCIerrors Displays the VCI lookup error cells.

ClockScale Lists the factor by which the clock is being scaled for traffic policing.

MulticastMode If the switch contains any Series A or B network modules, the switch operates in mixed (non-extended) mode. Otherwise, the switch uses the added Series C functionality and operates in extended mode. For complete information about non-extended and extended mode, please see Appendix A, Hardware Maintenance Procedures, in the ATM Switch User's Manual.

B.3.3 Displaying the Board Topology

This command displays the SPANS topology of the ATM network of which this switch is a part. All SPANS-NNI links appear in the topology. Enter the following parameters:

```
localhost::configuration board> topology

B Source           Ipaddress       Destination      Ipaddress       Capacity Age
1 f21a013a.08.0    198.29.22.46    f21a013a.09.0    198.29.22.46       1544    0
1 f21a013a.09.0    198.29.22.46    f21a013a.08.0    198.29.22.46       1544    0
```

The fields in this display are defined as follows:

B Indicates the number of the board (switch fabric).

Source Lists the source SPANS address of the link.

IpAddress Shows the IP address mapping to the source SPANS address, if it is known.

Destination Indicates the destination SPANS address of the link.

IpAddress Shows the IP address mapping to the destination SPANS address, if it is known.

Capacity Displays the link capacity in Kbps. A negative value in this field indicates that the link has gone down, but it has not timed out yet.

Age Shows the age of the link. A value of -1 indicates that the link is no longer present.

B.4 IP Configuration Commands

These commands let the user change the IP configuration. The user can display the list of available subcommands by typing `ip ?` at the `configuration` level.

```
localhost::configuration> ip ?
address         admin           broadcast       forwarding
mask            route>          show
```

B.4.1 Configuring the IP Address

This command allows the user to configure the IP address of each of the switch's IP interfaces. Enter the following parameters:

NOTE: On a new switch, the `ie0`, `asx0`, `qaa0`, `qaa1`, `qaa2`, `qaa3` interfaces are NOT configured. An IP address must be configured for at least one of the interfaces to allow IP access to the switch, which, in turn, enables SNMP access. By setting the IP address of the `asx0` interface or one of the `qaa` interfaces, in-band (over ATM) access to the switch control processor (SCP) is enabled. By setting the IP address of the `ie0` interface, out-of-band access to the SCP is enabled.

NOTE: The IP addresses must be configured individually on each SCP on an SFCS-1000.

```
localhost::configuration ip> address <interface> <address>
```

interface Indicates the name of the IP interface to be managed. Valid interfaces are: `ie0` (the Ethernet interface), `asx0` (the switch's SPANS interface), `qaa0`, `qaa1`, `qaa2`, `qaa3` (the Classical IP interfaces), and `lo0` (the switch's localhost interface that allows AMI to run).

address Indicates the IP address for this interface. The state of the interface must be `up` before setting the address.

This can be changed using the `configuration ip admin` command.

B.4.2 Configuring the IP State

This command allows the user to change the state of the IP interface to up or down. Enter the following parameters:

```
localhost::configuration ip> admin <interface> (up|down)
```

interface Indicates the name of the IP interface to be managed. Valid interfaces are: `ie0`, `asx0`, `qaa0`, `qaa1`, `qaa2`, and `qaa3`.

up|down Entering `up` changes the state of the designated interface to up. Entering `down` changes the state of the designated interface to down.

NOTE: The switch's localhost interface, `lo0`, must always be `up` to allow AMI to run on the switch.

B.4.3 Configuring the IP Broadcast Address

This command allows the user to modify the broadcast address for an IP interface. Enter the following parameters:

```
localhost::configuration ip> broadcast <interface> (0|1)
```

interface Indicates the name of the IP interface. Valid interfaces are: `ie0` , `asx0`, `qaa0`, `qaa1`, `qaa2`, and `qaa3`.

0|1 Indicates the IP broadcast type for this interface. This is the host portion of the IP address that is used for routing. Entering `1` causes the host portion of the IP address to be set to all 1s. Entering `0` causes the host portion of the IP address to be set to all 0s.

B.4.4 Configuring IP Forwarding

This command allows the user to turn IP forwarding on or off. If IP forwarding is turned off, the switch will not forward (i.e., route) IP packets from one IP interface to another IP interface. It is generally not necessary to turn IP forwarding off, except for security reasons. Enter the following parameters:

```
localhost::configuration ip> forwarding (on|off)
```

on|off Using on turns IP forwarding on. Using off turns IP forwarding off.

B.4.5 Configuring the IP Subnet Mask

This command allows the user to modify the IP subnet mask. Enter the following parameters:

```
localhost::configuration ip> mask <interface> <mask>
```

interface Indicates the name of the IP interface. Valid interfaces are: ie0, asx0, qaa0, qaa1, qaa2, and qaa3.

mask Indicates the subnet mask for this IP interface. It should be entered in dotted decimal notation (e.g., 255.255.255.0).

B.4.6 Configuring IP Routes

This command allows the user to add a static IP route to the local IP routing table, delete a static IP route from the local IP routing table, or list the current static IP routes in the local IP routing table. The user can display the list of available subcommands by typing `route ?` at the `ip` level.

```
localhost::configuration ip> route ?
new             delete          show
```

B.4.6.1 Adding an IP Route

This command allows the user to create an IP route. Enter the following parameters:

```
localhost::configuration ip route> new (default|<destination-ipaddress>) <gateway> [<metric>]
[(host | net)]
```

default — This parameter must be specified to create a default route.

destination-ipaddress — Indicates the destination IP network number.

gateway — Indicates the gateway address to the destination IP network number.

metric — Indicates the number of hops to the destination IP network. The default value of `1` is used if no value is entered by the user. If `1` is specified, the route is created with the RTF_GATEWAY flag set.

host | net — Using `host` indicates that this is a host-specific route with the destination being a specific node's IP address. Using `net` indicates that this is a network-specific route with the destination being a network IP address. The default value of `net` is used if no value is entered by the user or if a value is entered incorrectly (e.g., if a typo is made).

B.4.6.2 Deleting an IP Route

This command lets the user delete an IP route. Enter the following parameters:

```
localhost::configuration ip route> delete (default|<destination-ipaddress>) <gateway>
```

default A default must be specified to delete a default route.

destination-ipaddress Indicates the destination IP network number.

gateway Indicates the gateway address to the destination IP network number.

B.4.6.3 Showing the IP Routes

This command lets the user display the current IP routes. Enter the following parameters:

```
localhost:: configuration ip route> show

Destination      Gateway             Metric          Interface       Flags
default          169.144.48.1         1                le0             G
169.144.48.0     169.144.48.21        0                le0
169.144.60.0     169.144.60.21        0                asx0
169.144.64.0     169.144.64.21        0                qaa0
169.144.200.0    169.144.200.21       0                el0
169.144.204.0    169.144.204.21       0                el1
169.144.205.0    169.144.205.21       0                el2
169.144.206.0    169.144.206.21       0                el3
```

The fields in this display are defined as follows:

Destination Displays the destination IP network.

Gateway Shows the gateway address to the destination IP network number.

Metric Lists the number of hops to the destination IP network. The default is 1.

Interface Displays the local IP interface used to get to the destination IP network.

Flags Lists H if the route is host-specific (created with the RTF_HOST flag set). Lists G if the route is network-specific (created with the RTF_GATEWAY flag set).

B.4.7 Displaying the IP Interface Configuration

This command allows the user to display information about the configuration of the IP interfaces. Enter the following parameters:

```
localhost::configuration ip> show

interface  state       address         netmask         broadcast
ie0        up          169.144.48.21   255.255.255.0   169.144.48.0
asx0       up          169.144.60.21   255.255.255.0   169.144.60.0
qaa0       up          169.144.64.21   255.255.255.0   169.144.64.0
qaa1       down
qaa2       down
qaa3       down
lo0        up          127.0.0.1       255.0.0.0       127.0.0.0
el0        up          169.144.200.21  255.255.255.0   169.144.200.0
el1        up          169.144.204.21  255.255.255.0   169.144.204.0
el2        up          169.144.205.21  255.255.255.0   169.144.205.0
el3        up          169.144.206.21  255.255.255.0   169.144.206.0

IP Forwarding State: forwarding
```

The fields in this display are defined as follows:

interface Indicates the name of the IP interface.

state Lists the administrative state of the IP interface.

address Displays the IP address of the IP interface.

netmask Shows the netmask address of the IP interface.

broadcast Indicates the broadcast address of the IP interface.

The user may also designate a single interface to be displayed by entering show and the specific interface name at the prompt as follows:

```
localhost::configuration ip> show ie0

interface  state       address         netmask         broadcast
ie0        up          169.144.48.21   255.255.255.0   169.144.48.0

IP Forwarding State: forwarding
```

The fields in this display are defined in the same manner as those listed above in the example for showing the configuration of all of the IP interfaces.

B.5 LAN Emulation Configuration Commands

These commands allow the user to configure LAN Emulation (LANE) on a switch. The user can display the list of available subcommands by typing `lane ?` at the `configuration` level.

```
localhost::configuration> lane ?
  bus>              lec>              lecs>             les>
```

B.5.1 Broadcast and Unknown Server (BUS) Configuration Commands

These commands allow the user to configure a Broadcast and Unknown Server (BUS) for an ELAN. The user can reach this level by entering `bus` at the `lane` level. Enter the following parameters and type `?` to list the various `bus` commands:

```
localhost::configuration lane> bus ?
  admin             delete            new               show
```

B.5.1.1 Configuring the BUS Administrative Status

This command lets the user change the administrative status of a BUS to up (start a BUS service) or down (stop a BUS service). Enter the following parameters:

```
localhost::configuration lane bus> admin <BUS index> (up | down)
```

BUS index — Indicates the unique, positive integer index of the BUS that is dynamically assigned by AMI when a BUS is created to identify this service from any other service in the same class. The index can be found under the `Index` field when you enter the `conf lane bus show` command.

up | down — Entering `up` changes the administrative status of the designated BUS index to up. Entering `down` changes the administrative status of the designated BUS index to down. The default is `up`.

B.5.1.2 Deleting a BUS

This command allows the user to delete a specified BUS. Enter the following parameters:

```
localhost::configuration lane bus> delete <BUS index>
```

BUS index Indicates the unique, positive integer index of the BUS that is to be deleted. This number is dynamically assigned by AMI when a BUS is created. The index can be found under the `Index` field when you enter the `conf lane bus show` command.

B.5.1.3 Creating a BUS

This command lets the user create a BUS for an ELAN. Enter the following parameters:

```
localhost::configuration lane bus> new <BUS Selector byte (HEX)> <BUS name>
```

BUS Selector byte (HEX) Indicates the 20th byte of the ATM address of the switch that is to run a BUS service (entered in hexadecimal format). Use the `conf atmarp getnsap` command to display the entire ATM address.

BUS name Indicates the name for this BUS. It helps the user to remember which ELAN this BUS services.

B.5.1.4 Displaying BUS Information

This command lets the user display the current BUS information in one of two ways. To display information about every BUS that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane bus> show
 Index  AdminStatus  OperStatus   Selector  Name
     1  up           up              0x0b  marketing
```

The fields in this display are defined as follows:

Index Shows the unique, positive integer index that identifies this BUS. It is dynamically assigned by AMI when the BUS is created.

AdminStatus Reflects any changes that the user has made to the status of the BUS. Up means the user started the BUS. Down means the user stopped the BUS.

OperStatus Reflects the actual current status of the BUS. Up means the BUS is currently active. Down means the BUS is currently inactive.

Selector Displays the selector byte portion (20th byte) of the ATM address of the host or switch that is the BUS in hexadecimal format.

Name Lists the name of the ELAN that this BUS services.

NOTE: When you change the administrative status of a BUS from down to up, it takes a few seconds for the operational change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the information above to show the AdminStatus as up, but the OperStatus as down. If you show the information again after two or three seconds, the change will have taken place and be reflected here.

To display information about a particular BUS that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane bus> show [<BUS index>]
```

For example, to display information about the BUS with an index number of 1, enter the following parameters:

```
localhost::configuration lane bus> show 1


localhost::configuration lane bus> show
 Index  AdminStatus  OperStatus   Selector  Name
     1  up           up               0x0b  marketing
```

The fields in this display are defined in the same manner as those listed above in the example showing every BUS configured on the switch.

B.5.2 LAN Emulation Client (LEC) Configuration Commands

These commands let the user configure the LAN Emulation Client (LEC). The user can display the list of available subcommands by typing `lec ?` at the `lane` level.

```
localhost::configuration lane> lec ?
  admin            arp>             delete           default>
  new              show
```

B.5.2.1 Configuring the LEC Administrative Status

This command lets the user change the administrative status of a LEC to up (start a LEC) or down (stop a LEC). Enter the following parameters:

```
localhost::configuration lane lec> admin <LEC index> (up | down)
```

LEC index Indicates the unique, positive integer index of the LEC that is dynamically assigned by AMI when a LEC is created to identify this LEC from any others in the same ELAN. The index can be found under the `Index` field when you enter the `conf lane lec show` command.

up | down Entering `up` starts this LEC. Entering `down` stops this existing LEC. The default is `up`.

B.5.2.2 Configuring LANE ARP Commands

These commands let the user configure the LANE ARP cache. The user can reach this level by entering `arp` at the `lec` level. Enter the following parameters and type `?` to list the various ARP commands:

```
localhost::configuration lane lec> arp ?
  delete           show
```

B.5.2.2.1 Deleting LANE ARP Cache Information

This command allows the user to remove an ARP entry from the LANE ARP cache or to delete the contents of the LANE ARP cache. Enter the following parameters:

```
localhost::configuration lane lec arp> delete all | <MAC address>
```

all Indicates that all of the entries are to be flushed from the LANE ARP cache.

<MAC address> Indicates the specific entry that is to be flushed from the LANE ARP cache.

B.5.2.2.2 Displaying LANE ARP Cache Information

This command displays the current LANE ARP cache. The MAC address-to-ATM address mapping information for each LEC is stored here. Enter the following parameters:

```
localhost::configuration lane lec arp> show [(advanced)]
```

By entering show without the advanced option, the basic LANE ARP cache information is displayed as follows:

```
localhost::configuration lane lec arp> show
 MacAddress   AtmAddress                                          ELAN
 0020481a00d5 0x47.0005.80.ffe100.0000.f21a.00d5.0020481a00d5.0b eng-net
```

By entering show with the advanced option, more LANE ARP cache information, including the VPI/VCI combination and any flags associated with this entry, is displayed as follows:

```
localhost::configuration lane lec arp> show advanced
 MacAddress   AtmAddress                                          ELAN
 0020481a00d5 0x47.0005.80.ffe100.0000.f21a.00d5.0020481a00d5.0b eng-net
        vpi=0, vci=82, flags=valid
```

If the LANE ARP cache is empty, then the following message is displayed.

```
No LANE ARP entries are available.
```

B.5.2.3 Deleting a LEC

This command lets the user delete a LEC from an ELAN. Enter the following parameters:

```
localhost::configuration lane lec> delete <LEC index>
```

LEC index Indicates the unique, integer index of the LEC that is to be deleted. This number is dynamically assigned by the switch when a LEC is created and can be found under the `Index` field when you enter the `conf lane lec show` command.

B.5.2.4 LEC Default Configuration Mode Commands

These commands allow the user to set or to display the default LEC configuration mode. The user can show the list of available subcommands by typing `default ?` at the `lec` level.

```
localhost::configuration lane lec> default ?
  mode              show
```

B.5.2.4.1 Setting the Default LEC Configuration Mode

This command lets the user set the default mode for configuring all of the ELANs that may be created on this switch.

> **NOTE:** If you chose manual mode, you must specify the LECS address of the machine that will be used as the LECS. If you choose automatic mode, then the "well-known" LECS address is used.

Enter the following parameters:

```
localhost::configuration lane lec default> mode (manual | automatic) [<LECS address>]
    LECS address is required for manual mode.
```

manual | automatic Using manual means that the LECS address specified here is used as the LECS address. Using automatic means that the "well-known" LECS address (47.0079.00.000000.0000.0000.0000.00A03E00000 1.00) as defined by the ATM Forum's LAN Emulation standard is used as the LECS address. The default is automatic.

LECS address Indicates the ATM address of the LECS to be used instead of the "well-known" LECS.

B.5.2.4.2 Displaying the Default LEC Configuration Mode

This command lets the user show whether the default LEC configuration mode is manual (using a LECS other than the one at the "well-known" address) or automatic (using the LECS at the "well-known" address). Enter the following parameters:

```
localhost::configuration lane lec default> show
LEC Default configuration mode: automatic
```

B.5.2.5 Creating a LEC

This command lets the user create a LEC (join an ELAN). When a LEC is created, a corresponding el interface is created. The interface name (el0, el1, el2, etc.) is assigned based on the selector byte entered when the LEC is created. The list of current el interfaces can be displayed using the conf lane lec show command or the conf ip show command. Enter the following parameters:

NOTE: This AMI command only allows you to create an instance of a LEC on a switch. To create an instance of a LEC on a host, you must use the *ForeRunner* VLAN Manager or use a *ForeRunner* adapter. Please refer to the respective User's Manual for instructions.

```
localhost::conf lane lec> new <LEC Selector byte (HEX)> <ELAN name> [(automatic | manual)]
manual mode options: [-lecs <LECS address>] or [-les <LES address>]
Use ELAN name "default" to join default ELAN
```

LEC Selector byte (HEX) — Indicates the 20th byte of the ATM address of the LEC (entered in hexadecimal format). Use the conf atmarp getnsap command to display the entire ATM address.

ELAN name — Indicates the name of the ELAN that this LEC is joining. If a failover mechanism has been set up in the LECS configuration file, use the name of the failover LEC (e.g., eng-net|0). For more information, refer to Chapter 3 of this manual.

automatic | manual — Indicates the configuration mode that is used when this LEC joins the ELAN. Using automatic means that the the LEC attempts to contact the LECS using the "well-known" address. Using manual and the -lecs option means that the LEC attempts to contact the LECS using the LECS address you specified here. Using manual and the -les option means that the LEC bypasses the LECS and directly contacts the LES address specified here. The default is automatic.

-lecs <LECS address> — Indicates the LECS address to use instead of the "well-known" LECS address.

-les <LES address> — Indicates the LES address to use for this ELAN.

B.5.2.6 Displaying LEC Information

This command lets the user display the current LEC information in one of two ways. To display information about all of the LECs that are currently configured on the switch, enter the following parameters:

```
localhost::configuration lane lec> show
        Admin   Oper
 Index  Status  Status  Sel   Mode        MACaddress    IfName    ELAN
    1  up      up      0x00  automatic   000000000000  FAILOVER  eng-net|0
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.a0
    2  up      up      0x01  automatic   000000000000  FAILOVER  eng-net|1
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.a1
    3  up      up      0x02  automatic   00204815096b  el0       eng-net|2
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f215.096b.00204815096b.a2
    4  up      up      0x03  automatic   000000000000  FAILOVER  sw-net|0
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f215.096b.00204815096b.a0
    5  up      joining 0x04  automatic   0a204815096b  el1       sw-net|1
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f215.096b.00204815096b.a1
    6  up      joining 0x05  automatic   000000000000  FAILOVER  sw-net|2
       LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
       LES :0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.a2
```

The fields in this display are defined as follows:

Index — Shows the unique, integer index that identifies this LEC. It is dynamically assigned by the switch when the LEC is created.

AdminStatus — Reflects any changes that the user has made to the status of the LEC. `Up` means the user started the LEC. `Down` means the user stopped the LEC.

OperStatus — Reflects the actual current status of the LEC. `Up` means the LEC is currently active. `Down` means the LEC is currently inactive. `Joining` means that the LEC is in the process of registering with the ELAN.

Sel — Shows the selector byte portion (20th byte) of the ATM address of the LEC in hexadecimal format.

Mode Shows the configuration mode that is used when a LEC joins the ELAN. Automatic means that the "well-known" LECS address and the default LES are used. Manual means that the specified LECS or LES address is used. The default is automatic.

MACaddress Shows the Ethernet MAC address for this LEC.

IfName Shows the el interface name of this LEC. If it is part of a failover sequence, the the el interface name (e.g., el1) is displayed for the active el interface and FAILOVER is displayed for the others.

ELAN Shows the name of the ELAN to which this LEC belongs.

NOTE: When you change the administrative status from down to up, it takes a few seconds for the operational change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the information above to show the AdminStatus as up, but the OperStatus as down. If you show the information again after two or three seconds, the change has taken place and is reflected here.

To display information about a particular LEC that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane lec> show [<LEC index>]
```

For example, to display information about the LEC with an index number of 3, enter the following parameters:

```
localhost::configuration lane lec> show 3
        Admin   Oper
 Index  Status  Status  Sel   Mode        MACaddress    IfName    ELAN
    3  up      joining  0x02  automatic   00204815096b  el0      eng-net|2
        LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
        LES :0x47.0005.80.ffe100.0000.f215.096b.00204815096b.a2
```

The fields in this display are defined in the same manner as those listed above in the example showing all of the LECs configured on the switch.

B.5.3 LAN Emulation Configuration Server (LECS) Commands

These commands allow the user to configure the LAN Emulation Configuration Server (LECS). The user can display the list of available subcommands by typing `lecs ?` at the `lane` level.

```
localhost::configuration lane> lecs ?
  admin            delete           new              show
  get
```

B.5.3.1 Configuring the LECS Administrative Status

This command lets the user change the administrative status of the LECS to up (start a LECS service) or down (stop a LECS service). Enter the following parameters:

```
localhost::configuration lane lecs> admin <LECS index> (up | down)
```

LECS index Indicates the unique, positive integer index of the LECS that is dynamically assigned by AMI when a LECS is created to identify this service from any other service in the same class. The index can be found under the `Index` field when you enter the `conf lane lecs show` command.

up | down Entering `up` changes the administrative status of the designated LECS index to up. Entering `down` changes the administrative status of the designated LECS index to down. The default is `up`.

B.5.3.2 Deleting a LECS

This command lets the user delete (stop) a specified LECS service. Enter the following parameters:

```
localhost::configuration lane lecs> delete <LECS index>
```

LECS index Indicates the unique, positive integer index of the LECS that is to be deleted. This number is dynamically assigned by the switch when a LECS is created and can be found under the `Index` field when you enter the `conf lane lecs show` command.

B.5.3.3 Creating a LECS

This command lets the user create (start) a LECS service. Enter the following parameters:

```
localhost::configuration lane lecs> new <LECS Selector byte (HEX)> [-db <LECS database file>]
    [-default <LES atm address>]
```

LECS Selector byte (HEX) — Indicates the 20th byte of the ATM address of the host or switch that is to run a LECS service (entered in hexadecimal format). Use the `conf atmarp get-nsap` command to display the entire ATM address.

-db <LECS database file> — Indicates the full path to the location and name of the LECS database file. The default file for a 9A000, SFCS-200BX, SFCS-200WG, and an SFCS-1000 is `lecs.cfg`. For information about configuring this file, refer to Chapter 3 of this manual.

-default <LES atm address> — Indicates a default LES address to use in case the LECS configuration file is inaccessible.

NOTE: Only one LECS can be created per switch because the LECS listens for requests on a well-known address. However, more than one LEC, BUS, or LES may be created on a switch.

B.5.3.4 Displaying LECS Information

This command lets the user display the current LECS information in one of two ways. To display information about all of the LECs that are currently configured on the switch, enter the following parameters:

```
localhost::configuration lane lecs> show
 Index  AdminStatus  OperStatus   Selector  Database
     1  up           up              0x0c  lecs.cfg
        Default LES :00000000000000000000000000000000000000000000
```

The fields in this display are defined as follows:

Index Shows the unique, integer index of the LECS that is dynamically assigned by the switch when the LECS is created.

AdminStatus Reflects any changes that the user has made to the status of the LECS. `Up` means the user started the LECS. `Down` means the user stopped the LECS.

OperStatus Reflects the actual current status of the LECS. `Up` means the LECS is currently active. `Down` means the LECS is currently inactive.

Selector Displays the selector byte portion (20th byte) of the ATM address of the host or switch that is the LECS in hexadecimal format.

Database Shows the full path to the location and name of the LECS database file. The default file for a 9A000, SFCS-200BX, an SFCS-200WG, and an SFCS-1000 is `lecs.cfg`.

Default LES Shows the default LES address to use in case the LECS configuration file is inaccessible. If one has not been specified, shows all zeros for the address.

NOTE: When you change the administrative status of a LECS down to up, it takes a few seconds for the operational change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the information above to show the AdminStatus as up, but the OperStatus as down. If you show the information again after two or three seconds, the change will have taken place and will be reflected here.

To display information about a particular LECS that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane lecs> show [<LECS index>]
```

For example, to display information about the LECS with an index number of 1, enter the following parameters:

```
localhost::configuration lane lecs> show 1
 Index  AdminStatus  OperStatus   Selector  Database
     1  up           up              0x0c  lecs.cfg
        Default LES :00000000000000000000000000000000000000000000
```

The fields in this display are defined in the same manner as those listed above in the example for all of the LECS configured on the switch.

B.5.3.5 Getting the LECS Configuration File

This command lets the user download the LECS configuration file. Enter the following parameters:

```
localhost::configuration lane lecs> get <host>:<remotefile> [<localfile>]
```

host Indicates the name of the host from which the LECS database file is to be retrieved.

remotefile Indicates the name of the LECS database file that is to be retrieved.

localfile Indicates the name of the file where the retrieved LECS database file is to be stored.

NOTE: The default local file for a 9A000, SFCS-200BX, SFCS-200WG, and an SFCS-1000 is `lecs.cfg`.

NOTE: On a 9A000, SFCS-200BX, an SFCS-200WG, and an SFCS-1000, this file will be retrieved via tftp.

NOTE: For information about configuring this file, refer to Chapter 3 of this manual.

B.5.4 LAN Emulation Server (LES) Configuration Commands

These commands allow the user to configure the LAN Emulation Server (LES). The user can display the list of available subcommands by typing `les ?` at the `lane` level.

```
localhost::configuration lane> les ?
  admin            delete           new              show
```

B.5.4.1 Configuring the LES Administrative Status

This command lets the user change the administrative status of the LES to up (start a LES service) or down (stop a LES service). Enter the following parameters:

```
localhost::configuration lane les> admin <LES index> (up | down)
```

LES index Indicates the unique, positive integer index of the LES that is dynamically assigned by AMI when a LES is created to identify this service from any other service in the same class. The index can be found under the `Index` field when you enter the `conf lane les show` command.

up | down Entering `up` changes the administrative status of the designated LES index to up. Entering `down` changes the administrative status of the designated LES index to down. The default is `up`.

B.5.4.2 Deleting a LES

This command lets the user delete a specified LES. Enter the following parameters:

```
localhost::configuration lane les> delete <LES index>
```

LES index Indicates the unique, positive integer index of the LES that is to be deleted. This number is dynamically assigned by the switch when a LES is created and can be found under the `Index` field when you enter the `conf lane les show` command.

B.5.4.3 Creating a LES

This command lets the user create a new LES. Enter the following parameters:

```
localhost::conf  lane  les>  new <LES Selector byte (HEX)> <BUS ATM address> <LES name> \
[(colocated_bus)]
```

LES Selector byte (HEX) Indicates the 20th byte of the ATM address of the host or switch that is to run a LES service (entered in hexadecimal format). Use the `conf atmarp getnsap` command to display the entire ATM address.

BUS ATM address Indicates the ATM address of the BUS associated with this LES. If you are starting a colocated BUS, then you only need to enter the BUS selector byte instead of the full address. If the BUS is running on a different switch, you must enter the full address. Use the `conf atmarp getnsap` command to display the entire ATM address.

LES name Indicates the user-defined name associated with this LES to help the user to remember to what ELAN this LES belongs.

colocated_bus If `colocated_bus` is entered when the LES is created, indicates that the LES and BUS services for a particular ELAN are running on the same switch. It is recommended that the `colocated_bus` option be used when creating the LES.

NOTE: By using the `colocated_bus` option, you are creating a LES and BUS using a single AMI command. There is no need to create a BUS separately.

B.5.4.4 Displaying LES Information

This command lets the user display the current LES information in four different ways. To display information about every LES that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane les> show

Index  AdminStatus  OperStatus   Selector  ELAN
    1  up           up              0x0a  eng-net
      BUS:0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.0b (Co-Located)
    2  up           up              0x1a  eng-net2
      BUS:0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1b (Co-Located)
```

The fields in this display are defined as follows:

Index Shows the unique, integer index of the LES that is dynamically assigned by the switch when the LES is created.

AdminStatus Reflects any changes that the user has made to the status of the LES. Up means the user started the LES. Down means the user stopped the LES.

OperStatus Reflects the actual current status of the LES. Up means the LES is currently active. Down means the LES is currently inactive.

Selector Displays the selector byte portion (20th byte) of the ATM address of the host or switch that is the LES in hexadecimal format.

ELAN Shows the name of the ELAN that this LES services.

NOTE: When you change the administrative status of a LES up to down, it takes a few seconds for the operational change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the information above to show the AdminStatus as up, but the OperStatus as down. If you show the information again after two or three seconds, the change will have taken place and is reflected here.

To display information about a particular LES that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane les> show [<LES index>]
```

For example, to display information about the LES with an index number of 2, enter the following parameters:

```
localhost::configuration lane les> show 2


Index  AdminStatus  OperStatus   Selector  ELAN
    2  up           up              0x1a  eng-net2
      BUS:0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1b (Co-Located)
```

The fields in this display are defined in the same manner as those listed above in the example showing every LES configured on the switch.

To display advanced information about every LES that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane les> show advanced

ELAN Name: “eng-net”
  LES: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.0a
  BUS: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.0b
  Non-proxy Control Distribute VCC: 0.80
      Proxy Control Distribute VCC: -.-
LEC #1 at 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.00 (non-proxy)
  00:20:48:10:2a:78 -> 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.00
  Control Direct VCC: 0.77
LEC #4 at 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.10 (non-proxy)
  00:20:48:1a:00:d0 -> 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.10
  Control Direct VCC: 0.98
LEC #5 at 47.0005.80.ffe100.0000.f215.096b.00204815096b.10 (non-proxy)
  00:20:48:15:09:6b -> 47.0005.80.ffe100.0000.f215.096b.00204815096b.10
  Control Direct VCC: 0.140
LEC #6 at 47.0005.80.ffe100.0000.f215.096b.0020481013f6.00 (non-proxy)
  00:20:48:10:13:f6 -> 47.0005.80.ffe100.0000.f215.096b.0020481013f6.00
  Control Direct VCC: 0.144
```

```
ELAN Name: “eng-net2”
  LES: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1a
  BUS: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1b
  Non-proxy Control Distribute VCC: 0.79
      Proxy Control Distribute VCC: -.-
LEC #1 at 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.02 (non-proxy)
  06:20:48:10:2a:78 -> 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.02
  Control Direct VCC: 0.76
LEC #4 at 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.20 (non-proxy)
  06:20:48:1a:00:d0 -> 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.20
  Control Direct VCC: 0.101
LEC #5 at 47.0005.80.ffe100.0000.f215.096b.00204815096b.20 (non-proxy)
  06:20:48:15:09:6b -> 47.0005.80.ffe100.0000.f215.096b.00204815096b.20
  Control Direct VCC: 0.141
LEC #6 at 47.0005.80.ffe100.0000.f215.096b.0020481013f6.02 (non-proxy)
  06:20:48:10:13:f6 -> 47.0005.80.ffe100.0000.f215.096b.0020481013f6.02
  Control Direct VCC: 0.145
```

The fields in this display are defined as follows:

ELAN Name	Shows the names of any ELANs associated with this switch.
LES	Displays the ATM address of the LES that services this particular ELAN.
BUS	Displays the ATM address of the BUS that services this particular ELAN.
Non-proxy Control Distribute VCC	Identifies the point-to-multipoint connection that the LES maintains to all of the non-proxy LECs that it services.
Proxy Control Distribute VCC	Identifies the point-to-multipoint connection that the LES maintains to all of the proxy LECs that it services. This entry is blank if no proxy LECs have joined this ELAN.
LEC	Shows the LEC ATM address, shows the MAC-to-ATM address mapping for each LEC in this ELAN, and shows whether or not this LEC is a proxy.
Control Direct VCC	Shows the point-to-point connection that the LES maintains to this particular LEC.

To display advanced information about a particular LES that is currently configured on the switch, enter the following parameters:

```
localhost::configuration lane les> show [<LES index>] [(advanced)]
```

For example, to display advanced information about the LES with an index number of 2, enter the following parameters:

```
localhost::configuration lane les> show 2 advanced

ELAN Name: “eng-net2”
  LES: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1a
  BUS: 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.1b
  Non-proxy Control Distribute VCC: 0.79
      Proxy Control Distribute VCC: -.-
LEC #1 at 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.02 (non-proxy)
  06:20:48:10:2a:78 -> 47.0005.80.ffe100.0000.f21a.00d0.002048102a78.02
  Control Direct VCC: 0.76
LEC #4 at 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.20 (non-proxy)
  06:20:48:1a:00:d0 -> 47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.20
  Control Direct VCC: 0.101
LEC #5 at 47.0005.80.ffe100.0000.f215.096b.00204815096b.20 (non-proxy)
  06:20:48:15:09:6b -> 47.0005.80.ffe100.0000.f215.096b.00204815096b.20
  Control Direct VCC: 0.141
LEC #6 at 47.0005.80.ffe100.0000.f215.096b.0020481013f6.02 (non-proxy)
  06:20:48:10:13:f6 -> 47.0005.80.ffe100.0000.f215.096b.0020481013f6.02
  Control Direct VCC: 0.145
```

The fields in this display are defined in the same manner as those listed above in the example showing advanced information for every LES configured on the switch.

B.6 Network Module Configuration Commands

These commands let the user configure the network modules. The list of available subcommands can be displayed by typing `module ?` at the `configuration` level.

```
localhost::configuration> module ?
show            timing>          traffic>
```

B.6.1 Displaying Network Module Configuration Information

This command allows the user to display information about the configuration of the network modules that are currently installed in a particular switch board. Enter the following parameters:

```
localhost::configuration module> show
Module  Speed  Ports  Uptime    Type
4B      622.0      1  0d:17:58  NM-C-OC12c/STM4c-TIMING-SMIR-128KB-1PT (Rev. 1.0)
4E     2560.0      4  0d:17:58  NM-ASX1000-BPSwitch Backplane Netmod (Rev. 1551712.1)
```

The fields in this display are defined as follows:

Module Shows each network module that is currently in the switch board. The 4 means that it is the switch fabric in slot 4. B indicates the position of the physical network module in the switch fabric. E is the intra-fabric ports to the other switch fabrics in the chassis.

Speed Shows the speed in Mbps of the ports on the network module.

Ports Shows the number of ports on the network module.

Uptime Displays the length of time that the network module has been in its current state.

Type NM stands for network module. C is the series of network module. OC12c shows what kind of network module it is. TIMING means that it supports distributed timing. SMIR means singlemode intermediate reach. 128KB is the shared memory size. 1PT is the number of ports on the module. The E represents the intra-fabric ports to the other switch fabrics in the chassis. Rev 1.0 is the hardware revision.

B.6.2 Configuring Distributed Timing on a Network Module

These commands enable the user to configure or to display information about the timing source of the network modules.

> **NOTE:** These commands apply only to FORE Systems' Series C network modules that have distributed timing support.

Enter the following parameters to show the commands for network module timing:

```
localhost::configuration module> timing ?
show           internalclock          exportclock           importclock
```

B.6.2.1 Displaying the Network Module Timing Source

This command allows the user to display information about the timing source of the network modules.

> **NOTE:** These commands apply only to FORE Systems' Series C network modules that have distributed timing support.

Enter the following parameters:

```
localhost::configuration module timing> show

                 Import                          Export
Module  Internal  Primary   Secondary Current   Primary   Secondary Current
1C      export    N/A       N/A       N/A       crystal   crystal   primary
```

The fields in this display are defined as follows:

Module Shows each distributed timing network module that is currently installed in the switch. The 1 means that it is the first switch fabric. The C indicates the position of the network module in the switch fabric.

Internal Shows whether the export clock or the import clock is being used as the internal clock for this network module.

Import Primary Defines the preferred import timing source for the specified network module. This can be designated as the export clock of a different network module on the same switch fabric.

Import Secondary Defines the backup import timing source to be used for the specified network module in the event that the import primary source is unavailable. This can be designated as the export clock of a different network module on the same switch fabric.

Import Current Lists the current import timing source for the specified network module. In the event that the primary source is unavailable, it reads `secondary`.

Export Primary Defines the preferred export timing source for the specified network module. This clock source may be either recovered from one of the ports on this network module or this clock source may be the crystal oscillator on this network module.

Export Secondary Defines the backup export timing source to be used for the specified network module in the event that the export primary source is unavailable. This clock source may be either recovered from one of the ports on this network module or this clock source may be the crystal oscillator on this network module.

Export Current Lists the current export timing source for the specified network module. Normally, it shows `primary`. In the event that the primary source is unavailable, it reads `secondary`. If the export primary and export secondary timing sources are both unavailable (i.e., both sources were defined as ports), the on-board crystal is used as the export timing source. The default setting is `crystal`.

B.6.2.2 Configuring the Internal Clock Timing of a Network Module

Each network module installed in a switch has its own internal clock as a timing source. This timing source can be either the export clock or it may be the import clock. This command allows the user to configure the internal clock on a specified network module.

> **NOTE:** These commands apply only to FORE Systems' Series C network modules that have distributed timing support.

Enter the following parameters:

```
localhost::configuration module timing> internalclock <module> (export | import)
```

module Indicates the specific distributed timing network module to be configured.

export Indicates the export clock is being used by this network module as the internal clock. This clock source may be either recovered from one of the ports on this network module or this clock source may be the crystal oscillator on this network module.

import Indicates the import clock is being used by this network module as the internal clock. This clock source is being taken from the export clock of a different network module on the same switch fabric.

B.6.2.3 Configuring the Export Clock Timing of a Network Module

This command allows the user to configure the timing source for the export clock on a specified network module.

> **NOTE:** These commands apply only to FORE Systems' Series C network modules that have distributed timing support.

Enter the following parameters:

```
localhost::configuration module timing> exportclock <module> (primary | secondary)
(<mod-port#> | crystal | none)
```

module — Indicates the specific distributed timing network module to be configured.

primary | secondary — Using primary designates this as the preferred export timing source for this network module. Using secondary designates this as the backup export timing source to be used for this network module in the event that the primary source is unavailable.

mod-port# | crystal | none — Specifies the source of the export clock for this network module. Using mod-port# means that the timing source is recovered externally from this specific port on a distributed timing network module. Using crystal means that the timing is derived internally from the crystal oscillator on this network module. None is only available on the TP25 network module. Using none allows the TP25 network module to disable transmitting sync pulses.

For example, to configure the fourth port on network module 1B as the secondary source of timing for the export clock, enter the following parameters:

```
localhost::configuration module timing> exportclock 1B secondary 4
```

Note that for the final parameter, mod-port# , the user only needs to enter the port number (e.g., 4), not the module and port number (e.g., 1B4), since the second parameter is the module number.

B.6.2.4 Configuring the Import Clock Timing of a Network Module

This command allows the user to configure the timing source for the import clock on a specified network module.

> **NOTE:** These commands apply only to FORE Systems' Series C network modules that have distributed timing support.

Enter the following parameters:

```
localhost::configuration module timing> importclock <module> (primary | secondary) (A | B | C | D)
```

module	Indicates the specific distributed timing network module to be configured.
primary | **secondary**	Using primary designates this as the preferred import timing source for this network module. Using secondary designates this as the backup import source to be used for this network module in the event that the primary source is unavailable.
A | **B** | **C** | **D**	Indicates the import timing source will be taken from the export clock from another (A - bottom left, B - bottom right, C - top left, or D - top right) network module on this switch fabric.

B.6.3 Configuring Traffic on a Network Module

These commands enable the user to configure or to display information about the traffic on the network modules.

NOTE: These commands apply to FORE Systems' Series C network modules only.

Enter the following parameters to show the commands for network module traffic:

```
localhost::configuration module> traffic ?
epd             models          setmodel        show
fifoblock
```

B.6.3.1 Setting Early Packet Discard on a Network Module

This command lets the user set a threshold value for AAL5 Early Packet Discard (EPD) on a specified network module. This is the static threshold (in terms of a specified number of cells) at which EPD is activated. AAL5 frames that arrive when the shared buffer is over this threshold are discarded in whole.

NOTE: This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration module traffic> epd <module> <number of cells>
```

module Indicates the network module to be configured for Early Packet Discard.

number of cells Indicates the EPD threshold to be set. By default, this value is set to 90% of the shared buffer size.

B.6.3.2 Displaying Traffic Models for a Network Module

This command allows the user to display the different types of traffic memory models on a network module.

NOTE: This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration module traffic> models

 Model Unicasts   Mcasts McastOuts    Cells  Name
 1         4096      512      1024     2048  small memory default
 2         6144      512      2048    12288  big memory default
 3        11264     1024      4096    10240  big memory, more connections
 4         6144     4096      8192     8192  big memory, more multicast
 5         2048     2048     32768     8192  big memory, VP traffic shaping
 6         3072      128      1024    13312  big memory, fewer connections
```

The fields in this display are defined as follows:

Model Lists the memory model for this configuration.

Unicasts Shows the maximum number of unicast connections supported for this model.

Mcasts Displays the number of input multicast connections to support.

McastOuts Lists the number of multicast outputs to support.

Cells Shows the total amount of memory that is available for cells.

Name Displays the identifier for this memory model.

B.6.3.3 Setting Traffic Models on a Network Module

This command lets the user select one of the traffic memory models for a specific network module.

NOTE: This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration module traffic> setmodel <module> <model>
```

module Indicates which network module is to be configured.

model Indicates the predefined memory model to be used for this network module. The various models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections.

NOTE: The switch software must be restarted for this command to take effect.

B.6.3.4 Displaying Traffic on a Network Module

This command enables the user to display traffic model information on the network modules.

> **NOTE:** This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration module traffic> show
Module Model Unicasts Mcasts McastOuts  Cells Shared  Used   EPD FIFOblock
1B         1       10      0         0      0   1232      0   1109 normal
```

The fields in this display are defined as follows:

Module Shows the network module that has been configured. The number indicates the switch fabric and the letter indicates the position of the network module in the switch fabric.

Model Indicates the memory model used for this network module.

Unicasts Displays the current number of active unicast connections on this network module.

Mcasts Shows the current number of multicast inputs on this network module.

McastOuts Lists the current number of active multicast outputs on this network module.

Cells Displays the current number of active cells on this network module.

Shared Indicates the amount of shared memory allocated to this network module.

Used Shows the amount of shared memory currently in use on this network module.

EPD Lists the threshold, specified in cells, for Early Packet Discard on this network module.

FIFOblock Displays whether or not FIFO blocking is enabled on this network module.

B.6.3.5 Setting FIFO Blocking on a Network Module

This command lets the user set FIFO blocking on a specific network module when the queues are full.

> **NOTE:** This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration module traffic> fifoblock <module> (normal | enabled)
```

module Indicates which network module is to be configured.

normal | enabled `Normal` indicates that the network module passes traffic normally. `Enabled` means that FIFO blocking takes place on the network module when the queues are full. The default state is `normal`.

B.7 NSAP Configuration Commands

These commands allow the user to create, delete, and display NSAP static routes; to create, delete, and display NSAP prefixes; and to display ILMI registered NSAP addresses. The user can display the list of available subcommands by typing `nsap ?` at the `configuration` level.

```
localhost::configuration> nsap ?

route>          prefix>          ilmi>
```

B.7.1 NSAP Route Configuration Commands

These commands let the user create, delete, and display NSAP static routes. The user can display the list of available subcommands by typing `route ?` at the `nsap` level.

```
localhost::configuration nsap> route ?
  delete          new             show
```

B.7.1.1 Deleting an NSAP Route

This command enables the user to remove an existing NSAP static route. Enter the following parameters:

```
localhost::configuration nsap route> delete <NSAP> <mask> <port> <vpi>
```

NSAP Indicates the complete 20-byte NSAP route address in hexadecimal format.

mask Indicates the bit mask indicating number of high-order bits to use for routing purposes. The default mask for the route to the host is 152 and the default mask for the route to the switch is 104.

port Specifies the port on which this NSAP route is to be deleted.

vpi Specifies the virtual path on which this NSAP static route is to be deleted.

B.7.1.2 Creating an NSAP Route

This command allows the user to create an NSAP static route. Enter the following parameters:

```
localhost::configuration nsap route> new <NSAP> <mask> <port> <vpi> [-cost <cost>] [-cbr_cap
<cbr_cap>] [-vbr_cap <vbr_cap>] [abr] [epd]
```

NSAP Indicates the complete 20-byte NSAP route address in hexadecimal format.

mask Indicates the bit mask indicating number of high-order bits to use for routing purposes. The default mask for a static route to a host is 152 and the default mask for a static route to another switch is 104.

port Specifies the port through which this NSAP route can be reached.

vpi Specifies the UNI 3.0 signalling path through which this NSAP route can be reached.

-cost <cost> Indicates the routing metric for this link. There is a cost for each link in a route. The sum of these link costs determines the overall cost of a route. To expedite traffic on a route, the user should try to minimize the overall cost of a route. For a critical route, then, choose a small cost value. For a lesser important route, choose a higher cost value. The default value is 100.

-cbr_cap <cbr_cap> Indicates the maximum CBR (Constant Bit Rate) capacity allowed for any single connection on this route. This number is limited by the actual CBR capacity available on the output link specified for this route.

-vbr_cap <vbr_cap> Indicates the maximum VBR (Variable Bit Rate) capacity allowed for any single connection on this route. This number is limited by the actual VBR capacity available on the output link specified for this route.

abr This is an optional parameter. The `abr` parameter should be used for links that support ABR traffic. The `abr` parameter should be not be used for links that do not support ABR traffic.

epd This is an optional parameter. The `epd` parameter should be used for links that support Early Packet Discard. The `epd` parameter should not be used for links that do not support Early Packet Discard.

The following is an example of how to create an NSAP static route:

```
localhost::configuration nsap route> new 0x47.0005.80.ffe100.0000.f215.11f2.002048100464.00 152
1c2 0 -cost 200 -cbr_cap 20000 -vbr_cap 30000 abr
```

B.7.1.3 Displaying NSAP Routes

This command lets the user display the current NSAP static routes. Enter the following parameters:

```
localhost::configuration nsap route> show
```

```
localhost::configuration nsap route> show
NSAP-address                              Mask Port VPI Cost CBR   VBR   FLAGS
                                                            Mbs   Mbs
47000580ffe1000000f21511f200204810046400 152  1C2  0   200  20.0  30.0  A
47000580ffe1000000f21511f20020481ee00000 144  1C3  0   100  70.0  60.0  AE
47000580ffe1000000f21511f20020481ff00000 144  1C1  0   100  INF   INF   AE
47000580ffe1000000f21511f20020481ff12300 152  1C3  0   100  INF   INF
```

The fields in this display have the following meaning:

NSAP-address Shows all current NSAP static routes on the switch.

Mask Lists the bit mask indicating number of high-order bits to use for routing purposes. The default mask for a static route to a host is 152 and the default mask for a static route to another switch is 104.

Port Lists the port number on which the NSAP route exists. The 1 indicates that it is the first switch fabric. The letter C indicates the position of the network module in the switch. The 1, 2, 3 indicate the specific port number on the network module.

VPI Shows the number of the virtual path on which the NSAP static route exists.

Cost Shows the routing metric for this link. There is a cost for each link in a route. The sum of these link costs determines the overall cost of a route. To expedite traffic on a route, the user should try to minimize the overall cost of a route. A small cost value is assigned to a critical route, while a higher cost value is assigned to a lesser important route. The default value is 100.

CBR Displays the maximum CBR capacity allowed for any single connection on this route. INF means that user did not specify a value for this parameter when the route was created, so the value defaults to the capacity available on the outgoing link.

VBR Displays the maximum VBR capacity allowed for any single connection on this route. INF means that user did not specify a value for this parameter when the route was created, so the value defaults to the capacity available on the outgoing link.

FLAGS A means this link supports ABR traffic. E means that this route supports Early Packet Discard (EPD). AE means this link supports both ABR traffic and EPD. No flags indicate that neither EPD nor ABR traffic are supported on this link.

B.7.2 NSAP Prefix Configuration Commands

These commands enable the user to delete an NSAP prefix, create an NSAP prefix, and display NSAP prefix information. The user can display the list of available subcommands by typing prefix ? at the nsap level.

```
localhost::configuration nsap> prefix ?
  delete           new              show
```

B.7.2.1 Deleting an NSAP Prefix

This command lets the user remove an existing NSAP prefix. Enter the following parameters:

```
localhost::config nsap prefix> delete <port> <vpi> <prefix>
```

port Indicates the port number on which the NSAP prefix is to be deleted.

vpi Indicates the number of the virtual path on which the NSAP prefix is to be deleted.

prefix Indicates the NSAP prefix for this entry.

B.7.2.2 Creating an NSAP Prefix

This command lets the user create an NSAP prefix. Enter the following parameters:

```
localhost::config nsap prefix> new <port> <vpi> <prefix>
```

port Indicates the port number on which the NSAP prefix is to be created.

vpi Indicates the number of the virtual path on which the NSAP prefix is to be created.

prefix Indicates the NSAP prefix for this entry.

NOTE: Because multiple prefixes are not supported on the same port, delete the old prefix before creating a new one.

B.7.2.3 Displaying NSAP Prefixes

This command enables the user to display the current list of NSAP prefixes. Enter the following parameters:

```
localhost::configuration nsap prefix> show

Port      VPI NSAP-Prefix
1A1         0 0x47.0005.80.ffe100.0000.f124.00de
1A2         0 0x47.0005.80.ffe100.0000.f124.00de
1A3         0 0x47.0005.80.ffe100.0000.f124.00de
1A4         0 0x47.0005.80.ffe100.0000.f124.00de
1B1         0 0x47.0005.80.ffe100.0000.f124.00de
1B2         0 0x47.0005.80.ffe100.0000.f124.00de
1B3         0 0x47.0005.80.ffe100.0000.f124.00de
1B4         0 0x47.0005.80.ffe100.0000.f124.00de
1B5         0 0x47.0005.80.ffe100.0000.f124.00de
1B6         0 0x47.0005.80.ffe100.0000.f124.00de
1CTL        0 0x47.0005.80.ffe100.0000.f124.00de
```

The fields in this display have the following meanings:

Port Indicates the port number of the network module that is currently installed in the switch. The 1 indicates that it is the first switch fabric. The letter A indicates that the network module is installed in the bottom left-hand slot in the switch. The letter B indicates that the network module is installed in the bottom right-hand slot in the switch. The 1, 2, 3, 4, 5, 6 indicate the specific port number on the network module. CTL indicates the logical control port.

VPI Indicates the number of the virtual path on which the NSAP prefix exists.

NSAP-Prefix Indicates the NSAP prefix for this port.

B.7.3 NSAP ILMI Configuration Command

This command enables the user to display the NSAP addresses of all of the ports on a switch fabric that have been registered via ILMI. ILMI address registration occurs between the switch and host. The switch sends the host its 13-byte NSAP prefix. If the host accepts the prefix, the host builds its own NSAP address by appending its 7-byte host specific part. The host returns the complete 20-byte NSAP address to the switch. If the switch accepts the address, the switch enters that information into its topology tables and all connections destined for that NSAP address are routed to that host. These registration messages are sent over the reserved channel VPI 0, VCI 16. The user can display the available subcommand by typing `ilmi ?` at the `nsap` level.

```
localhost::configuration nsap> ilmi ?
  show
```

B.7.3.1 Displaying NSAP Addresses Registered through ILMI

This command enables the user to display the NSAP addresses of all of the ports on a switch fabric that have been registered via ILMI. Enter the following parameters:

```
localhost::configuration nsap ilmi> show
Port  NsapAddress
1B1   47000580ffe1000000f215116f00204810046400
1D1   47000580ffe1000000f215116f00204810308600
```

Port Lists the port number for which an NSAP address has been registered via ILMI.

NsapAddress Shows the NSAP address that has been registered through ILMI for this port.

To display the NSAP addresses that have been registered via ILMI for a specific port, enter the following parameters:

```
localhost::configuration nsap ilmi> show [<port>]

localhost::configuration nsap ilmi> show 1B1
Port  NsapAddress
1B1   47000580ffe1000000f215116f00204810046400
```

B.8 Port Configuration Commands

These commands let the user manage the configuration of the various ports. The user can display the list of available subcommands by typing `port ?` at the `configuration` level.

```
localhost::configuration> port ?

  cdvt            e1>             e3>             ds1>
  tp25>           ds3>            j2>             policing
  sonet>          show            taxi>           traffic>
  vbrob           vbrbuffob
```

B.8.1 CDVT Port Configuration Commands

This command lets the user modify the input Cell Delay Variation Tolerance (CDVT) on a per-port basis. Enter the following parameters:

```
localhost::configuration port> cdvt <port> <us>
```

port Indicates the port number on which the CDVT is to be changed.

us Specifies the new value for the CDVT setting in microseconds.

B.8.2 E-3 Port Configuration Commands

These commands allow the user to modify various aspects of the configuration of the ports on an E-3 network module. The following E-3 commands are available only when an E-3 network module is installed in the switch fabric. The user can display the list of available subcommands by typing `e3 ?` at the `port` level.

```
localhost::configuration port> e3 ?
  emptycells       loopback         mode             scrambling
  show             timing
```

B.8.2.1 Configuring E-3 Port Empty Cells

Empty cells are cells that are sent as "filler" or place holders when there is no real data to send. By sending these cells, network modules that are synchronous in nature can keep an even flow of traffic moving so that distributed timing can work properly. This command lets the user change the type of cells sent as empty cells on an E-3 network module port. Enter the following parameters:

```
localhost::configuration port e3> emptycells <port> (idle | unassigned)
```

port Indicates the port number on which the type of empty cells is to be changed.

idle | unassigned Indicates the type of cells this port sends for filler when the port is not sending data. `idle` cells set the CLP bit = 1 and `unassigned` cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is `unassigned`. In general, it is not necessary to change this variable from the default setting.

B.8.2.2 Configuring the E-3 Port Loopback

This subcommand allows the user to designate the type of loopback on a port on an E-3 network module. Enter the following parameters:

```
localhost::configuration port e3> loopback <port> (cell | payload | diag | line | none)
```

port Indicates the port number on which the loopback mode is to be changed.

cell Choosing `cell` loopback means that the E-3 stream received from the network is unframed into ATM cells. The cells are then reframed and transmitted back to the network.

payload Selecting `payload` loopback indicates that the E-3 stream received from the network has the E-3 overhead bits re-inserted and is retransmitted to the network.

diag Choosing `diagnostic` loopback connects the receiver to the transmitter. The E-3 stream transmitted by the switch to a port is looped back to the switch. The E-3 stream is still transmitted to the network, but the incoming E-3 stream is ignored.

line Choosing `line` loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted out to the line. Cells that are switched to this port are not sent over the line.

none Selecting `none` designates that no loopback will take place on the port. This is the default setting.

B.8.2.3 Configuring E-3 Port Mode

This command allows the user to change the method used for cell delineation on an E-3 network module port. Enter the following parameters:

```
localhost::configuration port e3> mode <port> (plcp | hcs)
```

port Indicates the port number on which the type of framing is to be changed.

plcp | hcs Using `plcp` means the port uses PLCP (Physical Layer Convergence Protocol) framing (G.751) for cell delineation. Using `hcs` means the port uses HCS (Header Check Sequence) based framing (G.832) for cell delineation. The default is `hcs`.

B.8.2.4 Configuring E-3 Port Scrambling

This subcommand allows the user to change the scrambling mode on a port on an E-3 network module. Enter the following parameters:

```
localhost::configuration port e3> scrambling <port> (on | off)
```

port Indicates the port number on which the scrambling mode is to be changed.

on | off Using `on` indicates that cell payload scrambling is enabled on this port. Using `off` means that cell payload scrambling is disabled on this port. The default setting is `off`.

NOTE: The scrambling mode should be set to the same status on both the transmitting side and the receiving side.

B.8.2.5 Showing the E-3 Port Configuration

This command allows the user to display current information about an E-3 network module. Enter the following parameters:

```
localhost::configuration port e3> show

Port Carrier Status Mode Loopback ClockSource  Scrambling IdleCells
1D1  no      0x72   hcs  none     internal     off        unassigned
1D2  no      0x72   hcs  none     internal     off        unassigned
1D3  no      0x88   hcs  none     internal     off        unassigned
1D4  no      0x88   hcs  none     internal     off        unassigned
```

The fields in this display are defined as follows:

Port Indicates the port number of the network module that is currently installed in the switch. The 1 means that it is the first switch fabric. The letter D means that the E-3 network module is installed in the bottom right-hand slot in the switch. The 1, 2, 3, 4 specify the port number on the network module.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, yes is displayed. If a carrier has not been detected, no is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee that the signal is the proper frequency.

Status Indicates the E-3 line status of the port.

Mode Indicates the mode of operation for this port. Plcp means that the port uses PLCP framing for cell delineation. Hcs means that the port uses HCS cell delineation.

Loopback Indicates the loopback mode on this port. Can be one of the following: cell, payload, diagnostic, line, or none.

ClockSource For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is set to be derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

Scrambling On means that payload scrambling is enabled on this port. Off means that payload scrambling is disabled on this port.

IdleCells Shows the type of cells this port sends for filler when the port is not sending data. Idle cells set the CLP bit = 1 and unassigned cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is unassigned. In general, it is not necessary to change this variable from the default setting.

B.8.2.6 Configuring the E-3 Port Timing

This command allows the user to change the timing source on a port on an E-3 network module. Enter the following parameters:

```
localhost::configuration port e3> timing <port> (network | internal)
```

port Indicates the port number from which the timing is being derived.

network | internal Designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is set to be derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

B.8.3 TP25 Port Configuration Commands

These commands allow the user to modify various aspects of the configuration of the ports on a TP25 network module. The following tp25 commands are available only when a TP25 network module is installed in the switch fabric. The user can display the list of available subcommands by typing tp25 ? at the port level.

```
localhost::configuration port> tp25 ?
loopback       show
```

B.8.3.1 Configuring the TP25 Port Loopback

This subcommand allows the user to designate the type of loopback on a port on a TP25 network module. Enter the following parameters:

```
localhost::configuration port tp25> loopback <port> (line | none)
```

port Indicates the port number on which the loopback mode is to be changed.

line Choosing line loopback, also known as remote loopback, causes received data to be transferred to the upstream system as well as to be looped back to the transmitter.

none Selecting none designates that no loopback will take place on the port. This is the default setting.

B.8.3.2 Showing the TP25 Port Configuration

This command allows the user to display current information about a TP25 network module. Enter the following parameters:

```
localhost::configuration port tp25> show
Port Carrier Media Loopback RxTiming
1A1  no      UTP   none     Yes
1A2  no      UTP   none     Yes
1A3  no      UTP   none     Yes
1A4  no      UTP   none     Yes
1A5  no      UTP   none     Yes
1A6  no      UTP   none     Yes
1B1  yes     UTP   none     Yes
1B2  yes     UTP   none     Yes
1B3  no      UTP   none     Yes
1B4  yes     UTP   none     Yes
1B5  no      UTP   none     Yes
1B6  yes     UTP   none     Yes
```

The fields in this display are defined as follows:

Port Indicates the port number of the network module that is currently installed in the switch.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, yes is displayed. If a carrier has not been detected, no is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee that the signal is the proper frequency.

Media Displays what kind of physical medium is connected to the TP25 interface. UTP means that it is Unshielded Twisted Pair.

Loopback Indicates the loopback mode on this port. Can be either line or none.

RxTiming Indicates whether or not the port is receiving an 8kHz Timing Sync Marker. These markers can be used to derive an 8kHz signal that can be transmitted from all ports on the network module and back to the switch fabric (on switches that support timing features). No means that the port is not receiving sync pulses. Yes means that the port is receiving sync pulses.

B.8.4 DS-3 Port Configuration Commands

These commands allow the user to modify various aspects of the configuration of the ports on a DS-3 network module. The following DS-3 commands are available only when a DS-3 network module is installed in the switch fabric. The user can display the list of available subcommands by typing `ds3 ?` at the `port` level.

```
emptycells      framing         loopback        mode
scrambling      show            timing          length
```

B.8.4.1 Configuring DS-3 Port Empty Cells

Empty cells are cells that are sent as "filler" or place holders when there is no real data to send. By sending these cells, network modules that are synchronous in nature can keep an even flow of traffic moving so that distributed timing can work properly. This command lets the user change the type of cells sent as empty cells on a DS-3 network module port. Enter the following parameters:

```
localhost::configuration port ds3> emptycells <port> (idle | unassigned)
```

port Indicates the port number on which the type of empty cells is to be changed.

idle | unassigned Indicates the type of cells this port sends for filler when the port is not sending data. `idle` cells set the CLP bit = 1 and `unassigned` cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is `unassigned`. In general, it is not necessary to change this variable from the default setting.

B.8.4.2 Configuring the DS-3 Port Framing

This command allows the user to designate the mode to be used on a port on a DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> framing <port> (cchannel | cbit)
```

port Indicates the port number on which the framing is to be changed.

cchannel | cbit Indicates the type of framing for the port. Using cchannel (clearchannel) means that standard M23 framing is used on this port. The default is cbit (cbitparity).

B.8.4.3 Configuring the DS-3 Port Loopback

This command allows the user to designate the type of loopback on a port on a DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> loopback <port> (cell | payload | diag | line | none)
```

port Indicates the port number on which the loopback mode is to be changed.

cell Choosing cell loopback means that the DS-3 stream received from the network is unframed into ATM cells. The cells are reframed and transmitted back to the network.

payload Selecting payload loopback mean the DS-3 stream received from the network has the DS-3 overhead bits re-inserted and is retransmitted to the network.

diag Choosing diagnostic loopback connects the receiver to the transmitter. The DS-3 stream transmitted by the switch to a port is looped back to the switch. The DS-3 stream is still transmitted to the network, but the incoming DS-3 stream is ignored.

line Selecting line loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted out to the line. Cells that are switched to this port are not sent over the line.

none Selecting none designates that no loopback will take place. This is the default setting.

B.8.4.4 Configuring the DS-3 Port Mode

This command allows the user to change the type of framing on a port on a DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> mode <port> (plcp | hcs)
```

port Indicates the port number on which the type of framing is to be changed.

plcp | hcs Using `plcp` means the port uses PLCP (Physical Layer Convergence Protocol) framing for cell delineation. Using `hcs` means the port uses HCS (Header Check Sequence) based framing for cell delineation. The default is `hcs`.

B.8.4.5 Configuring the DS-3 Port Scrambling

This command allows the user to change the scrambling mode on a port on the DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> scrambling <port> (on | off)
```

port Indicates the port number on which the scrambling mode is to be changed.

on | off Using `on` means that cell payload scrambling is enabled on this port. Using `off` means that cell payload scrambling is disabled on this port. Only the payload of the ATM cells is scrambled.

NOTE: The scrambling mode should be set to the same status on both the transmitting side and the receiving side.

B.8.4.6 Showing the DS-3 Port Configuration

This command allows the user to display current information about all of the ports on the DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> show

Port Carrier Status Mode Framing Loopback ClockSrc Scrambling IdleCells Length
1A1  no      0x80   hcs  cbit    diag     internal off        unassigned Gt225
1A2  no      0x24   plcp cbit    none     internal off        unassigned Gt225
```

Port Indicates the port number of the network module that is currently installed in the switch.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, `yes` is displayed. If a carrier has not been detected, `no` is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee the signal is the proper frequency.

Status Indicates the DS-3 line status of the port.

Mode `Plcp` means the port uses PLCP (Physical Layer Convergence Protocol) framing for cell delineation. `Hcs` means the port uses HCS (Header Check Sequence) based framing for cell delineation.

Framing Indicates the type of framing used for the port. Can be `cchannel` or `cbit`. The default is `cbit`.

Loopback Indicates the loopback mode of the port. Can be one of the following: `cell`, `payload`, `diagnostic`, or `none`.

ClockSource For all network modules, `network` means that the timing for this port is derived externally from the incoming clock on this port. `Internal` means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

Scrambling `On` means payload scrambling is enabled on the port. `Off` means payload scrambling is disabled on the port.

IdleCells Shows the type of cells this port sends for filler when the port is not sending data. idle cells set the CLP bit = 1 and unassigned cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is unassigned. In general, it is not necessary to change this variable from the default setting.

B.8.4.7 Configuring the DS-3 Port Timing

This command allows the user to change the timing source on a port on a DS-3 network module. Enter the following parameters:

```
localhost::configuration port ds3> timing <port> (network | internal)
```

port Indicates the port number from which the timing is being derived.

network | internal Designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

B.8.4.8 Configuring DS-3 Port Line Length

This command lets the user change the line length of a DS-3 port to correspond to the physical cable attached to that port. Enter the following parameters:

```
localhost::configuration port ds3> length <port> (Lt225 | Gt225)
```

port Indicates the port number on which the line length is to be changed.

Lt225 Use if the physical cable is shorter than 225 ft.

Gt225 Use if the physical cable is greater than 225 ft. This is the default.

B.8.5 J-2 Port Configuration Commands

These commands let the user modify various aspects of the configuration of a J-2 network module. The following J-2 commands are available only when a J-2 network module is installed in the switch fabric. The user can display the list of available subcommands by typing j2 ? at the port level.

```
localhost::configuration port> j2 ?
line             loopback         show             timing
```

B.8.5.1 Configuring J-2 Port Line Length

This command enables the user to change the line length of a J-2 network module port. Enter the following parameters:

```
localhost::configuration port j2> line <port> (short | long)
```

port Indicates the port number on which the line length is being changed.

short | long Indicates the length of the physical cable attached to this port. If the line attached to the receive port has greater than 4 db of attenuation, then the line must be configured as long. If otherwise, then it must be configured as short. In general, if the cable is less than 20 feet, then configure the line as short.

B.8.5.2 Configuring J-2 Port Loopback

This command allows the user to configure the loopback mode on a J-2 port. Enter the following parameters:

```
localhost::configuration port j2> loopback <port> (line | diag | none)
```

port Indicates the port number on which the loopback mode is to be changed.

line Selecting `line` loopback connects the transmitter to the receiver. The data stream received from the line is retransmitted out to the line. Cells that are switched to this port are not sent over the line.

diag Choosing `diagnostic` loopback connects the receiver to the transmitter. The J-2 stream transmitted by the switch to a port is looped back to the switch. The stream is still transmitted over the cable, but the incoming stream is ignored.

none Selecting `none` designates that no loopback will take place. This is the default setting.

B.8.5.3 Showing J-2 Port Configuration

This command allows the user to display information about the configuration of the ports on a J-2 network module. Enter the following parameters:

```
localhost::configuration port j2> show

Port LineLength   Loopback     ClockSource
1D1  short        none         internal
1D2  short        none         internal
1D3  short        none         internal
1D4  short        none         internal
```

The fields in this display are defined as follows:

Port Indicates the port number of the network module.

LineLength Indicates the length of the physical cable that is attached to this port. Can be short or long.

Loopback Indicates the loopback mode of this port. Can be one of the following: none, line, or diagnostic.

ClockSource For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is set to be derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

B.8.5.4 Configuring J-2 Port Timing

This command allows the user to change the timing source on a port on a J-2 network module. Enter the following parameters:

```
localhost::configuration port j2> timing <port> (network | internal)
```

port Indicates the port number from which the timing is being derived.

network | internal Designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is set to be derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

B.8.6 Port Policing Configuration Command

This command lets the user decide whether or not traffic policing is enabled on a given port. Enter the following parameters:

```
localhost::configuration port> policing <port> (enabled | disabled)
```

port Indicates the specific port number on which traffic policing is to be enabled or disabled.

enabled | disabled Choosing enabled means that traffic policing will take place on this port. Choosing disabled means that traffic policing will not take place on this port.

B.8.7 SONET Port Configuration Commands

These commands allow the user to modify various aspects of the configuration of all of the ports on a SONET network module. The following SONET commands are available only when a SONET network module is installed in the switch fabric. The user can display the list of available subcommands by typing `sonet ?` at the `port` level.

NOTE: All 155 Mbps and 622 Mbps network modules use this same set of commands, regardless of whether they are singlemode, multimode, OC-3, OC-12, or UTP network modules.

```
localhost::configuration port> sonet ?
emptycells     loopback       mode           show
timing
```

B.8.7.1 Configuring SONET Port Empty Cells

Empty cells are cells that are sent as "filler" or place holders when there is no real data to send. By sending these cells, network modules that are synchronous in nature can keep an even flow of traffic moving so that distributed timing can work properly. This command lets the user change the type of cells sent as empty cells on a SONET network module port. Enter the following parameters:

```
localhost::configuration port sonet> emptycells <port> (idle | unassigned)
```

port Indicates the port number on which the type of empty cells is to be changed.

idle | unassigned Indicates the type of cells this port sends for filler when the port is not sending data. `idle` cells set the CLP bit = 1 and `unassigned` cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is `unassigned`. In general, it is not necessary to change this variable from the default setting.

B.8.7.2 Configuring SONET Port Loopback

This command enables the user to configure the type of loopback mode on a SONET port. Enter the following parameters:

```
localhost::configuration port sonet> loopback <port> (line | diag | none)
```

port Indicates the port number on which the loopback mode is to be changed.

line Selecting line loopback connects the transmitter to the receiver. The data stream received from the fiber is retransmitted out to the fiber. In line loopback, the port acts as if it were an optical repeater. Cells that are switched to this port are not sent over the fiber.

diag Choosing diagnostic loopback connects the receiver to the transmitter. The SONET stream transmitted by the fiber to a port is looped back to the fiber. The stream is still transmitted over the fiber, but the incoming stream is ignored.

none Selecting none designates that no loopback will take place. This is the default setting.

B.8.7.3 Configuring SONET Port Mode

This command lets the user designate the mode to be used on a SONET network module port. Enter the following parameters:

```
localhost::configuration port sonet> mode <port> (sonet|sdh)
```

port Indicates the port number on which the mode is to be changed.

sonet | sdh Indicates the mode of operation for this port. Can be sonet or sdh.

B.8.7.4 Showing the SONET Port Configuration

This command lets the user display information about the configuration of all of the ports on a SONET network module. Enter the following parameters:

```
localhost::configuration port sonet> show

Port Width  Carrier Status Line Mode  Loopback ClockSource  Scramb. IdleCells
1A1  sts3c  no      0x2    SMM  sonet none     internal     on      unassigned
1A2  sts3c  no      0x2    SMM  sonet none     internal     on      unassigned
1A3  sts3c  yes     0x1    SMM  sonet none     internal     on      unassigned
1A4  sts3c  no      0x2    SMM  sonet none     internal     on      unassigned
```

The fields in this display are defined as follows:

Port Indicates the port number of the network module that is currently installed in the switch fabric. The 1 means that it is the first switch fabric. The letter A means that the SONET network module is installed in the bottom left-hand slot in the switch fabric. The 1, 2, 3, 4 indicate the specific port number on the network module.

Width Indicates the type of the SONET path. sts3c is 155.52 Mbps and sts12c is 622.08 Mbps. The SDH transmission rate STM-1 is equivalent to SONET rate STS-3 and STM-4 is equivalent to STS-12. The user can not change this parameter.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, yes is displayed. If a carrier has not been detected, no is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee the signal is the proper frequency.

Status Indicates the SONET line status of the port.

Line Displays the line type for this interface. The line type for optical SONET signals may be SMSR (155 Mbps single-mode short reach), SMIR (622 Mbps single-mode intermediate reach), other (155 Mbps single-mode long reach), or MM (155 Mbps or 622 Mbps multi-mode) fiber. For electrical interfaces, the line type is UTP (155 Mbps Unshielded Twisted Pair).

Mode Indicates the mode of operation for this port. Can be sonet or sdh.

Loopback Indicates the loopback mode on this port. Can be one of the following: none, line, or diagnostic.

ClockSource For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

Scrambling On indicates that payload scrambling is enabled on this port. Off means that payload scrambling is disabled on this port.

IdleCells Indicates the type of cells this port sends for filler when the port is not sending data. Idle cells set the CLP bit = 1 and unassigned cells set the CLP bit = 0. Please refer to page 57 of the ATM Forum 3.0 Specification for more information. Idle = invalid cell pattern and unassigned = unassigned. The default is unassigned. In general, it is not necessary to change this variable from the default setting.

B.8.7.5 Configuring SONET Port Timing

This command allows the user to change the timing source on a port on a SONET network module.

> **NOTE:** This option is not available on an OC-12 network module because it always uses internal timing. To configure distributed timing on an OC-12 network module, use the configuration module timing commands.

Enter the following parameters:

```
localhost::configuration port sonet> timing <port> (network | internal)
```

port Indicates the port number from which the timing is being derived.

network | internal Designates the source of the transmit clock. For all network modules, network means that the timing for this port is derived externally from the incoming clock on this port. Internal means that the timing is derived from either the on-board crystal oscillator or is derived from a specific port number on a Series C network module that is defined by the user. If set to a specific port number, the transmit clock is derived from that port if available. If the port source becomes unavailable, the on-board crystal takes over as the transmit clock source.

B.8.8 Showing the Port Configuration

This command lets the user display port information about all of the ports on an individual switch fabric or about just a specified port. To show general information about all of the ports, enter the following parameters:

```
localhost::configuration port> show

Port Carrier   Mb/s iVPs iVCs    iBW oVPs oVCs    oBw Model
1C1  no       155.0    1    4  155.0    1    4  155.0 OC3
1C2  no       155.0    1    4  155.0    1    4  155.0 OC3
1C3  no       155.0    1    4  155.0    1    4  155.0 OC3
1C4  yes      155.0    1    4  155.0    1    4  155.0 OC3
1D1  no       155.0    1    2  155.0    1    3    0.0 OC3
1D2  no       155.0    1    5  155.0    1    4  155.0 OC3
1D3  no       155.0    1    4  155.0    1    4  155.0 OC3
1D4  no       155.0    1    4  155.0    1    4  155.0 OC3
1CTL yes       80.0    1   28    0.0    1   36    0.0 ASX
```

The fields in this display are defined as follows:

Port Indicates the port number of the network modules that are currently installed in the switch fabric. CTL indicates the control port, which is a logical (not physical) location where cells that are directed to the SCP itself are sent.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, yes is displayed. If a carrier has not been detected, no is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee the signal is the proper frequency.

Mb/s Displays the capacity of this port's link in Mbps.

iVPs Shows the number of incoming virtual paths on the port.

iVCs Indicates the number of incoming virtual channels on the port.

iBW Lists the amount of bandwidth (in Mbps) that has been reserved for the input link on the port.

oVPs Designates the number of outgoing virtual paths on the port.

oVCs Indicates the number of outgoing virtual channels on the port.

oBw Shows the amount of bandwidth (in Mbps) for the output link on the port.

Model Lists the type of network module. For the control port, lists ASX.

To show advanced information about all of the ports, enter the following parameters:

```
localhost::configuration port> show advanced

Port   CDVT  Policing VBROB BuffOB
1C1     250  enabled    100    100
1C2     250  enabled    100    100
1C3     250  enabled    100    100
1C4     250  enabled    100    100
1D1     250  enabled    100    100
1D2     250  enabled    100    100
1D3     250  enabled    100    100
1D4     250  enabled    100    100
1CTL   5000  enabled    100    100
```

The fields in this display are defined as follows:

Port Indicates the port number of the network modules that are currently installed in the switch fabric.

CDVT Shows the default value for the cell delay variation tolerance setting in microseconds.

Policing Displays whether or not traffic policing is enabled or disabled for this port.

VBROB Shows the bandwidth overbooking level configured on this port, specified as a percentage. Valid values are integers from 1 to 500. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

BuffOB Indicates the buffer overbooking level configured on this port, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

To list port information for just a specified port, (for example, port 1C1), enter the following parameters:

```
localhost::configuration port> show 1C1

Port Carrier   Mb/s iVPs iVCs    iBW oVPs oVCs    oBw Model
1C1  no       155.0    1    4  155.0    1    4  155.0 OC3
```

The fields in this display are defined in the same manner as those listed previously in the example for all of the ports on an individual switch fabric.

To list advanced port information for just a specified port, (for example, port 1C1), enter the following parameters:

```
localhost::configuration port> show 1C1 advanced

Port   CDVT  Policing VBROB BuffOB
1C1     250  enabled    100    100
```

The fields in this display are defined in the same manner as those listed previously in the example for all of the ports on an individual switch fabric.

B.8.9 TAXI Port Configuration Commands

These commands allow the user to modify various aspects of the configuration of all of the ports on a TAXI network module. The following TAXI commands are available only when a TAXI network module is installed in the switch fabric. The user can display the list of available subcommands by typing taxi ? at the port level.

```
localhost::configuration port> taxi ?
loopback         show
```

B.8.9.1 Configuring TAXI Port Loopback

This command allows the user to designate the type of loopback on a port on a TAXI network module. Enter the following parameters:

```
localhost::configuration port taxi> loopback <port> (diag | none)
```

port Indicates the port number on which the loopback mode is to be changed.

diag Choosing diagnostic loopback connects the receiver to the transmitter. The TAXI stream transmitted by the fiber to a port is looped back to the fiber. The stream is still transmitted over the fiber, but the incoming stream is ignored.

none Selecting none designates that no loopback will take place. This is the default setting.

B.8.9.2 Showing the TAXI Port Configuration

This command enables the user to display current information about all of the ports on a TAXI network module. Enter the following parameters:

```
localhost::configuration port taxi> show

Port Carrier    State      Obuf  Version Loopback   Model
1A1  yes       up        8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
1A2  no        down      8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
1A3  no        down      8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
1A4  no        down      8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
1A5  no        down      8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
1A6  yes       up        8192  C      none      NM-C-TAXI-100-ST-128KB-6PT
```

The fields in this display are defined as follows:

Port Indicates the port number of the network module that is currently installed in the switch fabric. The 4 means that it is the fourth switch fabric. The letter B indicates that the TAXI network module is installed in the bottom right-hand slot in the switch fabric. The 1, 2, 3, 4, 5, 6 refer to the specific port number on the network module.

Carrier Shows whether or not a carrier has been detected on the port. If a carrier has been detected, `yes` is displayed. If a carrier has not been detected, `no` is displayed. A carrier is detected when a signal is applied to the receive side of the port. It does not guarantee the signal is the proper frequency.

State Displays the current state of the port.

Obuf Lists the size in cells of the port's output buffer.

Version Shows the hardware version of the network module.

Loopback Indicates the loopback mode of this port. Can be either `none` or `diagnostic`.

Model Displays the type of network module. NM means network module, C means that it is series C hardware, TAXI means that it is a TAXI network module, 100 indicates its speed in Mbps, ST indicates the connector type, 128 KB is the shared memory size, and 6PT indicates that it is a six-port version.

B.8.10 Traffic Port Configuration Commands

These commands enable the user to configure various traffic features on an individual port on a Series C network module on the switch. The user can display the list of available subcommands by typing `traffic ?` at the `port` level.

```
localhost::configuration port> traffic ?
cdv             efci            qsize           show
```

B.8.10.1 Configuring Cell Delay Variation on a Port

On a Series C network module, there are two output queues that are 256 cells deep, by default, one for Constant Bit Rate (CBR) and one for Variable Bit Rate (VBR). The Cell Delay Variation (CDV) for CBR is calculated as the CBR cell queue depth (256 cells by default) multiplied by 1 cell time. The CDV for VBR is calculated as the VBR cell queue depth (256 cells by default) multiplied by 1 cell time plus the CBR CDV. The CDV also varies depending on the physical interface on which the link is running (e.g., a 155 Mbps connection versus a 45 Mbps connection). This command lets the user set the maximum CDV on a worst case basis that cells for a specified output port and priority (CBR or VBR) should incur. This number is used to determine the size of the buffers reserved for CBR and VBR traffic.

NOTE: This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration port traffic> cdv <port> (CBR | VBR) <CDV in microseconds>
```

port Indicates the port on which the CDV will be set.

CBR | VBR Indicates if the CDV is for output CBR traffic or for output VBR traffic.

CDV Specified in microseconds, indicates the cell delay variation, that an output cell experiences under the worst conditions.

NOTE: The switch control software must be restarted for this command to take effect.

B.8.10.2 Configuring EFCI on a Port

This command allows the user to designate the cell buffer threshold over which Available Bit Rate (ABR) cells and Unspecified Bit Rate (UBR) cells have their explicit forward congestion indicator (EFCI) code point set. When the EFCI code point is set, this signals congestion to downstream switch fabrics and to future ABR and UBR flow control mechanisms. Once this threshold is surpassed, EFCI continues to be set until the queue empties.

> **NOTE:** This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration port traffic> efci <port> (on | off) <threshold>
```

port Indicates the port on which the EFCI threshold will be set.

on | off On means the EFCI will be set when the threshold number is reached, signalling congestion. Off means the EFCI will be cleared when the threshold number is reached, indicating no congestion.

threshold Indicates the number of cells over which the ABR cells and the UBR cells will have EFCI set. The default value is 64 cells.

B.8.10.3 Configuring Port Queue Size

This command enables the user to designate the dedicated queue size for a given type of traffic on a specified port.

> **NOTE:** This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration port traffic> qsize <port> (CBR | VBR | ABR) <number of cells>
```

port Indicates the port on which the dedicated queue size will be set.

CBR | VBR | ABR Indicates for which type of traffic (CBR, VBR, or ABR) to set the dedicated queue size.

number of cells Indicates the queue size to be assigned to the traffic designated in the previous parameter. The default is 256 cells.

> **NOTE:** The switch control software must be restarted for this command to take effect.

B.8.10.4 Displaying Port Traffic and Priority Queue Information

This command lets the user display port traffic information and port priority queue information for all of the ports on all of the Series C network modules.

> **NOTE:** This command applies to FORE Systems' Series C network modules only.

Enter the following parameters:

```
localhost::configuration port traffic> show

Port configuration:
        CBR             VBR             ABR-UBR
 Port   Qsize     CDV   Qsize     CDV   Qsize EFCI-ON EFCI-OFF
 4A1      256     700     256    1400     256      64        0
 4A2      256     700     256    1400     256      64        0
 4A3      256     700     256    1400     256      64        0
 4A4      256     700     256    1400     256      64        0

Port priority queues:

               CLP       Qsize     Qsize
 Port Priority Threshold Dedicated Current      TxCells  LostCells
 4A1  ABR-UBR        256       256       0       110957          0
 4A1  VBR            256       256       0            0          0
 4A1  CBR            256       256       0            0          0
 4A2  ABR-UBR        256       256       0        17930          0
 4A2  VBR            256       256       0            0          0
 4A2  CBR            256       256       0            0          0
 4A3  ABR-UBR        256       256       0        17931          0
 4A3  VBR            256       256       0            0          0
 4A3  CBR            256       256       0            0          0
 4A4  ABR-UBR        256       256       0        17932          0
 4A4  VBR            256       256       0            0          0
 4A4  CBR            256       256       0            0          0
```

The fields in the port configuration display are defined as follows:

Port Shows the port that has been configured. The 4 indicates the fourth switch fabric, the letter indicates the position of the network module in the switch fabric, and the 1, 2, 3, or 4 indicates the specific port num-

ber.

CBR Qsize Displays the reserved queue size for CBR traffic. The default setting is 256 cells.

CBR CDV Shows the maximum cell delay variation for CBR traffic specified in microseconds.

VBR Qsize Displays the reserved queue size for VBR traffic. The default setting is 256 cells.

VBR CDV Shows the maximum cell delay variation for VBR traffic specified in microseconds.

ABR-UBR Qsize Displays the reserved queue size for ABR/UBR traffic. The default setting is 256 cells.

ABR-UBR EFCI-ON Shows the value at which the EFCI will be set (turned on) when the threshold number is reached, signalling congestion, for ABR and UBR traffic.

ABR-UBR EFCI-OFF Shows the value at which the EFCI will be cleared (turned off) when the threshold number is reached, indicating no congestion, for ABR and UBR traffic.

The fields in the port priority queues display are defined as follows:

Port Shows the port that has been configured. The 4 indicates the fourth switch fabric, the letter indicates the position of the network module in the switch fabric, and the 1, 2, 3, or 4 indicates the specific port number.

Priority Indicates the type of traffic for this port.

CLP Threshold Lists the value at which cells that have been tagged as non-conforming will be dropped for this priority and port. This parameter can not be changed.

Qsize Dedicated Displays the reserved queue size for the type of traffic specified in the Priority field. This parameter can not be changed.

Qsize Current Shows the amount of the reserved queue size that has already been used on this port for the specified type of traffic.

TxCells Lists the amount of traffic that has been transmitted on this port for the specified type of traffic.

LostCells Shows the number of cells dropped on this port for the specified type of traffic.

This command also lets the user show port traffic information and port priority queue information for an individual port on a Series C network module. Enter the following parameters:

```
localhost::configuration port traffic> show [<port>]

localhost::configuration port traffic> show 4A4

Port configuration:
        CBR             VBR             ABR-UBR
 Port   Qsize     CDV   Qsize     CDV   Qsize EFCI-ON EFCI-OFF
 4A4      256     700     256    1400     256      64        0

Port priority queues:

               CLP       Qsize     Qsize
 Port Priority Threshold Dedicated Current      TxCells  LostCells
 4A4  ABR-UBR        256       256       0        17932          0
 4A4  VBR            256       256       0            0          0
 4A4  CBR            256       256       0            0          0
```

The fields in these displays are defined in the same manner as those listed above in the example for displaying traffic on all of the ports on the Series C network modules.

B.8.11 VBROB Port Configuration Commands

This command is an advanced option that allows the user to set an output bandwidth overbooking level for VBR traffic on a particular port. Enter the following parameters:

```
localhost::configuration port> vbrob <port> <percent>
```

port Indicates the port number on which the bandwidth overbooking level for VBR traffic is to be changed.

percent Indicates the bandwidth overbooking level assigned to this port, specified as a percentage. Enter an integer value from 1 to 500. The default is 100, which indicates that no overbooking will occur. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

B.8.12 VBRBuffOB Port Configuration Commands

This command is an advanced option that allows the user to set an output buffer overbooking level for VBR traffic on a particular port. Enter the following parameters:

```
localhost::configuration port> vbrbuffob <port> <percent>
```

port Indicates the port number on which the buffer overbooking level for VBR traffic is to be changed.

percent Indicates the buffer overbooking level assigned to this path, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

B.9 Serial Port Configuration Commands

These commands let the user manage an RS-232 serial port. The user can display the list of available subcommands by typing rs232 ? at the configuration level.

```
localhost::configuration> rs232 ?
  show              speed
```

B.9.1 Displaying Serial Port Information

This command allows the user to show information about the configuration of the RS-232 serial port on the SCP. The user may type show at the rs232 configuration level to list the settings for the serial port. Enter the following parameters:

```
localhost::configuration rs232> show [(A|B)]
```

A | B Indicates the port letter of the serial port that is to be displayed.

```
localhost::configuration rs232> show

Port     Type        Speed   Flow     Bits   Stops      Parity
A        rs232        9600   none        8   one        none
```

The fields in this display are defined as follows:

Port Shows the physical port designation.

Type Indicates the signalling standard used.

Speed Displays the receive/transmit rate in bits per second.

Flow Indicates the type of flow control implemented on the given port.

Bits Shows the number of bit times in a single character.

Stops Lists the number of stop bits in a character frame.

Parity Displays the parity setting for the ports. Can be odd, even, mark, space, or none.

B.10 SNMP Configuration Commands

These commands enable the user to manage the SNMP communities and traps. The user can display the list of available subcommands by typing snmp ? at the configuration level.

```
localhost::configuration> snmp ?

community              trap>
```

B.10.1 Configuring the SNMP Community Access

This command lets the user modify the SNMP community access to AMI. Enter the following parameters:

```
localhost::configuration snmp> community (read | write) <community>
```

read | write Indicates the access level for this community. If set to read, the AMI session opens with read-only access. If set to write, the AMI session opens with read-write access.

community Indicates the community string associated with read or with write. The default community string associated with read is public. The default community string associated with write is private.

B.10.2 Configuring SNMP Traps

These commands help the user to manage SNMP traps. The user can display the list of available subcommands by typing trap ? at the snmp level.

```
localhost::configuration snmp> trap ?

delete         new            show
```

B.10.2.1 Deleting an SNMP Trap Entry

This command allows the user to delete an existing SNMP trap destination. Before deleting a trap that may need to be recreated later, show the list of current SNMP traps and either copy and save the screen or write down the trap destinations. You will also need to show the list of current SNMP traps in order to find the number of the trap to be deleted. Enter the following parameters to delete a trap entry:

```
localhost::configuration snmp trap> delete <trap>
```

trap Indicates the number of the trap destination in the list of current SNMP traps that is to be removed.

For example, to delete trap 198.29.31.130, first list the traps to find its number and copy the address in case you want to recreate it later:

```
localhost::configuration snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
5       198.29.31.130
```

Then enter the following parameters:

```
localhost::configuration snmp trap> delete 5
```

You can display the list again to verify that the trap has been deleted:

```
localhost::configuration snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
```

B.10.2.2 Creating an SNMP Trap Entry

This command allows the user to specify a host to which a switch can send SNMP traps. The SNMP traps supported by this switch are detailed in the FORE-Switch-MIB. Enter the IP address of the SNMP trap destination to be added. Repeat this for as many SNMP trap destinations as needed. Traps are active as soon as they are set. Enter the following parameters:

```
localhost::configuration snmp trap> new <ipaddress>
```

ipaddress Indicates the IP address of the trap destination to be created.

B.10.2.2.1 Displaying the SNMP Trap Entries

This command enables the user to list all of the current SNMP traps. The SNMP traps supported by this switch are detailed in the FORE-Switch-MIB. Enter the following parameters:

```
localhost::configuration snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
5       198.29.31.130
```

If no SNMP traps have been configured, the following message is displayed:

```
No trap information is available
```

B.11 SPANS Configuration Commands

These commands allow the user to manage SPANS (Simple Protocol for ATM Network Signalling), FORE Systems' pre-standard signalling protocol. The user can display the list of available subcommands by typing `spans ?` at the `configuration` level.

```
localhost::configuration> spans ?

delete          new             show
```

B.11.1 Deleting a SPANS Signalling Path

This command lets the user delete an existing SPANS signalling path. Enter the following parameters:

```
localhost::configuration spans> delete <port> <vpi>
```

port Indicates the port number on which the SPANS signalling path is to be deleted.

vpi This is the number of the SPANS path to be removed.

B.11.2 Creating a SPANS Signalling Path

This command allows the user to create a SPANS signalling path.

NOTE: Before a SPANS signalling path can be created on a given VPI, an originating and a terminating path must exist for that same VPI.

Enter the following parameters:

```
localhost::configuration spans> new <port> <vpi> [-cdvt <cdvt>] [(tag | drop)]
advanced options:
    [-sig <vci>] [-cls <vci>] [-aal (4 | 5)]
    [-sigbw <Kbps>] [-clsupc <index>]
    [-minvci <vci>] [-maxvci <vci>]
```

port Indicates the port number on which the SPANS signalling path is to be created.

vpi Indicates the number of the SPANS path that is to be created.

-cdvt <cdvt> Indicates the Cell Delay Variation Tolerance (CDVT) associated with the peak cell rates in microseconds.

(tag | drop) Using `tag` means that non-compliant cells will be tagged. Using `drop` means that non-compliant cells will be dropped.

The advanced options are as follows:

-sig <vci> Indicates the VCI to use for SPANS signalling messages. The default is `15`.

-cls <vci> Indicates the VCI to use for connectionless messages. The default is `14`.

-aal (4 | 5) Indicates the AAL type to use for this SPANS signalling path. The default is `4`.

-sigbw <Kbps> Indicates the amount of bandwidth to be reserved on the VCI for SPANS signalling messages.

-clsupc <index>	Indicates the integer index that refers to a specific traffic contract that is used to police the connectionless VCI. If no index is specified, then no traffic policing will take place on this VCI. It is assigned a UPC index of 0, and all traffic on this VCI is treated as UBR traffic. This is the default.
-minvci <vci>	Indicates the bottom number for the range of VCIs to be reserved for SPANS SVCs.
-maxvci <vci>	Indicates the top number for the range of VCIs to be reserved for SPANS SVCs.

B.11.3 Showing the SPANS Signalling Path Configuration

This command lets the user list an individual switch fabric's current SPANS signalling path information. Enter the following parameters:

```
localhost::configuration spans> show
 Port  VPI  State Type   CDVT  Action RemoteAddress
 1C1     0  down  uni     250  tag
 1C2     0  down  uni     250  tag
 1C3     0  down  uni     250  tag
 1C4     0  up    uni     250  tag    169.144.60.108
 1D1     0  down  uni     250  tag
 1D2     0  down  uni     250  tag
 1D3     0  down  uni     250  tag
 1D4     0  down  uni     250  tag
 1CTL    0  up    uni       0  tag    10.10.10.48
```

The fields in this display have the following meanings:

Port Lists the port number of the SPANS signalling path.

VPI Shows the path number of the SPANS signalling path.

State Shows the current state of the SPANS path.

Type Designates the type of connection on this SPANS path. If the type listed is uni, this is a SPANS user-to-network interface connection to a SPANS host. If the type listed is nni, then this is a SPANS network-to-network interface connection to another switch.

CDVT Shows the Cell Delay Variation Tolerance (CDVT), in microseconds.

Action Tag means that non-compliant cells are tagged. Drop means that non-compliant cells are discarded.

RemoteAddress Shows the IP address of the remote endstation, if it is available.

To show advanced SPANS signalling path information about all of the ports, enter the following parameters:

```
localhost::configuration spans> show advanced
 Port  VPI SigVCI CLSVCI AAL MinVCI MaxVCI SigBW CLSUPC
 1C1     0     15     14   4     32    511     0      0
 1C2     0     15     14   4     32    511     0      0
 1C3     0     15     14   4     32    511     0      0
 1C4     0     15     14   4     32    511     0      0
 1D1     0     15     14   4     32    511     0      0
 1D2     0     15     14   4     32    511     0      0
 1D3     0     15     14   4     32    511     0      0
 1D4     0     15     14   4     32    511     0      0
 1CTL    0     15     14   4     32    511     0      0
```

The fields in this display have the following meanings:

Port Lists the port number of the SPANS signalling path.

VPI Shows the path number of the SPANS signalling path.

SigVCI Indicates the virtual channel number used for SPANS messages on the SPANS path. The default is VCI 15.

ClsVCI Indicates the VCI used for connectionless messages.

AAL Lists the AAL type used for SPANS messages.

MinVCI Displays the bottom number for the range of VCIs to be reserved for SPANS SVCs.

MaxVCI Displays the top number for the range of VCIs to be reserved for SPANS SVCs.

SIGBW Lists the amount of bandwidth reserved on the VCI for SPANS signalling messages.

CLSUPC Shows the integer index that refers to a specific UPC contract used to police the connectionless VCI.

To list SPANS information for a specific port, (for example, port 1C1), enter the following parameters:

```
localhost::configuration spans> show 1C1
 Port  VPI  State Type   CDVT  Action RemoteAddress
 1C1     0  down  uni     250  tag
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying SPANS information on all of the ports on an individual switch fabric.

To list SPANS information for a specific port and path, (for example, port 1C1 and VPI 0), enter the following parameters:

```
localhost::configuration spans> show 1c1 0
 Port  VPI  State Type   CDVT  Action RemoteAddress
 1C1     0  down  uni     250  tag
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying SPANS information on all of the ports on an individual switch fabric.

To list advanced SPANS information for a specific port and path, (for example, port 1C1 and VPI 0), enter the following parameters:

```
localhost::configuration spans> show 1c1 0 advanced
 Port  VPI SigVCI CLSVCI AAL MinVCI MaxVCI SigBW CLSUPC
 1C1     0     15     14   4     32    511     0      0
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying advanced SPANS information on all of the ports on an individual switch fabric.

B.12 SPVC Configuration Commands

These commands allow the user to configure SPVCs (Smart Permanent Virtual Circuits). An SPVC is a connection that traverses multiple switch fabrics. An SPVC looks like a PVC at the local and remote endpoints with an SVC in the middle. SPVCs are more robust than PVCs. If a link carrying a PVC goes down, then the PVC goes down. If a link carrying an SPVC goes down and there is an alternate route, then the end switch fabrics of the SPVC automatically reroute the SPVC around the failed link.

The user can display the list of available subcommands by typing `spans ?` at the `configuration` level.

```
localhost::configuration> spvc ?
delete                new                   show
```

B.12.1 Deleting an SPVC

This command allows the user to delete an existing SPVC. There are two different ways to delete an SPVC. Enter the following parameters:

```
localhost::configuration spvc> delete <Local SPVC ID> [(source | destination | bidirectional)]
or
localhost::conf spvc> delete <port> <vpi> <vci> \ <dest-session> <dest-port> <dest-vpi> <dest-vci>\
[(source | destination | bidirectional)]
```

Local SPVC ID — Indicates the unique number that the SCP assigned to this SPVC when the SPVC was created.

source | destination | bidirectional — `Source` means the SPVC to be deleted is a unidirectional SPVC going from the local switch fabric to the remote switch fabric. `Destination` means the SPVC to be deleted is a unidirectional SPVC going from the remote switch fabric to the local switch fabric. `Bidirectional` means the pair of unidirectional SPVCs will be deleted. The default is `bidirectional`.

port — Indicates the port number on the local switch fabric on which this SPVC is to be deleted.

vpi — Indicates the virtual path number on the local switch fabric.

vci Indicates the virtual channel number on the local switch fabric.

dest-session Indicates the name of the remote switch session for this SPVC.

dest-port Indicates the port number on the remote switch fabric on which this SPVC is to be deleted.

dest-vpi Indicates the virtual path number on the remote switch fabric.

dest-vci Indicates the virtual channel number on the remote switch fabric.

source | destination | bidirectional Source means the SPVC to be deleted is a unidirectional SPVC going from the local switch fabric to the remote switch fabric. Destination means the SPVC to be deleted is a unidirectional SPVC going from the remote switch fabric to the local switch fabric. Bidirectional means the pair of unidirectional SPVCs will be deleted. The default is bidirectional.

B.12.2 Creating an SPVC

This command allows the user to create a new SPVC. Enter the following parameters:

```
localhost::conf spvc> new <port> <vpi> <vci> <dest-session> <dest-port> <dest-vpi> <dest-vci> \
[-peak <Kb/sec>] [(source | destination | bidirectional)]
```

port Indicates the port number on the local switch fabric on which this SPVC is to be created.

vpi Indicates the virtual path number on the local switch fabric.

vci Indicates the virtual channel number on the local switch fabric.

dest-session Indicates the name of the remote switch session.

dest-port Indicates the port number on the remote switch fabric on which this SPVC is to be created.

dest-vpi Indicates the virtual path number on the remote switch fabric.

dest-vci Indicates the virtual channel number on the remote switch fabric.

-peak <Kb/sec> Indicates the amount of peak bandwidth allocated for this SPVC, specified in kilobits per second. The default value is 0.

source | destination | bidirectional Source means a unidirectional SPVC going from the local switch fabric to the remote switch fabric will be created. Destination means a unidirectional SPVC going from the remote switch fabric to the local switch fabric will be created. Bidirectional means the pair of unidirectional SPVCs will be created. The default direction, if the user does not specify one, is bidirectional.

NOTE: To create a bidirectional SPVC, you must either specify bidirectional, or you must set up two unidirectional SPVCs with one going in each direction.

To create an SPVC, you need to configure the two ends concurrently on the two switch fabrics. Therefore, you first need to open an AMI session to the destination switch fabric by either using the SCP's IP address or its name, along with the SNMP read-write community string. The following example shows how to create a bidirectional SPVC from the local switch fabric (localhost) to a remote switch fabric (198.29.22.46). The user is logged into localhost.

```
localhost::> open 198.29.22.46 private

Opening a session for "198.29.22.46", please wait...
Connected to "198.29.22.46" (asx200bx).
198.29.22.46::> localhost

localhost::> configuration spvc new ?

usage: new <port> <vpi> <vci> <dest-session> <dest-port> <dest-vpi>
<dest-vci> \[-peak <Kb/sec>] [(source | destination | bidirectional)]

localhost::configuration spvc> new 1c1 0 49 198.29.22.46 1b1 0 50
```

B.12.3 Displaying SPVC Information

This command allows the user to display all of the SPVCs on an individual switch fabric. Enter the following parameters:

```
localhost::configuration spvc> show

Local                                    Remote
ID    Port VPI VCI     BW Direction      ID    Port VPI VCI Switch
35664 1C1    0  51    0.0 bidirectional  10427 1B1    0  52 198.29.22.46
65364 1C1    0  49    0.0 bidirectional  42591 1B1    0  50 198.29.22.46
```

The fields in this display are defined as follows:

Local ID Indicates the unique number that the local switch fabric's SCP assigned to this SPVC when the SPVC was created.

Local Port Indicates the port number on the local switch fabric.

Local VPI Indicates the virtual path number on the local switch fabric.

Local VCI Indicates the virtual channel number on the local switch fabric.

Local BW Indicates the amount of peak bandwidth allocated for this SPVC, specified in kilobits per second.

Remote ID Indicates the unique number that the remote switch fabric's SCP assigned to this SPVC when the SPVC was created.

Remote Port Indicates the port number on the remote switch fabric.

Remote VPI Indicates the virtual path number on the remote switch fabric.

Remote VCI Indicates the virtual channel number on the remote switch fabric.

Switch Indicates the IP address or name of the remote switch fabric's SCP.

B.13 Switch Configuration Commands

These commands allow the user to configure default settings for the switch. The user can display the list of available subcommands by typing switch ? at the configuration level.

```
localhost::configuration> switch ?
  name              pmpmaxvci         pmpminvci         show
```

B.13.1 Setting or Changing the Switch Name

This command enables the user to set or change the name of the switch. The switch name is shown on the front panel display LED. Enter the following parameters:

```
localhost::configuration switch> name <name>
```

name Indicates the new host name for the switch.

B.13.2 Setting the Minimum Number of Reserved VCIs for PMPs

This command lets the user set the minimum number for the range of VCIs that are reserved for point-to-multipoint connections. By using this command in conjunction with `conf switch pmpmaxvci`, you will allow a block of VCIs to be reserved for LAN Emulation point-to-multipoint use. This block of VCIs will be reserved on all paths and on all ports on this switch fabric. PVCs can be created on these VCIs, but no point-to-point connections may use these VCIs.

NOTE: This command is useful only when the switch is running in non-extended mode. For more information about non-extended mode, please see the section on Switch Board Configuration Commands found earlier in this appendix and the section on Network Module Replacement found in Appendix A of the Cabletron ATM Switch User's Manual.

Enter the following parameters:

```
localhost::configuration switch> pmpminvci <vci>
```

vci Indicates the bottom number for the range of VCIs to be reserved for point-to-multipoint connections. The default is 155.

B.13.3 Setting the Maximum Number of Reserved VCIs for PMPs

This command lets the user set the maximum number for the range of VCIs that are reserved for point-to-multipoint connections. By using this command in conjunction with `conf switch pmpminvci`, you will allow a block of VCIs to be reserved for LAN Emulation point-to-multipoint use. This block of VCIs will be reserved on all paths and on all ports on this switch fabric. PVCs can be created on these VCIs, but no point-to-point connections may use these VCIs.

NOTE: This command is useful only when the switch is running in non-extended mode. For more information about non-extended mode, please see the section on Switch Board Configuration Commands found earlier in this appendix and the section on Network Module Replacement found in Appendix A of the Cabletron ATM Switch User's Manual.

Enter the following parameters:

```
localhost::configuration switch> pmpmaxvci <vci>
```

vci Indicates the top number for the range of VCIs to be reserved for point-to-multipoint connections. The default is 255.

B.13.4 Displaying the Switch Configuration

This command lets the user display switch configuration information including the switch name, the type of switch, the hardware version, the software version, the number of maximum virtual paths, the number of maximum virtual channels, the SPANS address of the switch, and the range of the minimum and maximum number of reserved VCIs for point-to-multipoint connections. Enter the following parameters:

```
Switch 'fishtank', Type asx200bx, up 2 days 19:40
Hardware version 1.0, Software version S_ForeThought_4.0.0 (1.19)
Maximum Virtual Path Connections        32768
Maximum Virtual Channels                16384
SPANS address                           00000038f2150e87
PMP Minimum Reserved VCI                155
PMP Maximum Reserved VCI                255
```

NOTE: If the display reads 'ATM Switch' in the first line, this indicates that the switch name has not been set. Use the AMI command configuration switch name to assign a name.

B.14 System Configuration Commands

These commands let the user configure system message log features, configure the amount of time of non-activity after which an AMI session times out, and change the units for designating UPC contracts. The user can display the list of available subcommands by typing system ? at the configuration level.

```
localhost::configuration> system ?
  show              syslog>           timeout           units
```

B.14.1 Displaying System Information

This command lets the user display the amount of time of non-activity after which an AMI session will time out and display what type of units will be used when designating UPC contracts. The user can get to this level by entering show at the system level. Enter the following parameters:

```
localhost::configuration system> show
AMI Session Timeout                                    60
UPC Units                                              cps
```

AMI Session Timeout Shows the number of minutes of no activity after which an AMI session will time out and exit the user out of the session. The default is 5 minutes. A value of off means that the AMI session will not time out.

UPC Units Cps indicates that UPC contracts are being configured and displayed in cells per second. Kbps means that UPC contracts are being configured and displayed in kilobits per second. The default is cps.

B.14.2 System Log Configuration Commands

Syslog is a tool that can send system messages to be logged to a user-specified remote host. These commands let the user configure the address of the remote syslog host and whether or not these messages are sent to the console. Enter `syslog ?` at the `system` level to show the list of available `syslog` commands.

```
localhost::configuration system> syslog ?
  show              set               delete            console
```

B.14.2.1 Displaying the Address of the System Log Host

This command allows the user to display the address of the host to which the switch's system messages are logged. Enter the following parameters:

> **NOTE:** This command is only available on the local switch.

```
localhost::configuration system syslog> show
Remote Syslog Host: 169.144.1.216
Syslog Facility: daemon
```

If the host's address has never been set, or if it has been deleted and not set again, the following is shown:

```
localhost::configuration system syslog> show
No remote syslog host set. Syslog messages will not be sent.
Syslog Facility: daemon
```

B.14.2.2 Setting the Address of the System Log Host

This command sets the address of the host to which the switch's system messages are being logged. You may also opt to assign a specific facility name so that the remote syslog can automatically differentiate between switches with different facilities. Enter the following parameters:

NOTE: This command is only available on the local switch.

```
localhost::configuration system syslog> set <address> [<facility>]
```

address Indicates the IP address of the remote host to which the switch's system message logs are sent.

facility Indicates the user-assigned name that identifies the facility this switch sends syslogs on. The default is daemon. Other valid values are local0 through local7. The names may be assigned in any order.

For example, to have a host with the address 192.88.243.118 be the recipient of the system log messages and to designate this switch's facility name as local2, enter the following parameters:

```
localhost::configuration system syslog> set 192.88.243.118 local2
Remote Syslog Host: 192.88.243.118
Syslog Facility is now local2
```

To assign the facility for other switches in your network, you must log in to each one locally and set the facility. You must keep track of which facility was assigned to each switch. Then, when you view the contents of the syslog file, they can be separated according to facility when messages from different facilities are logged on the same remote host.

NOTE: If the switch panics, the panic file is automatically written to the syslog, provided that a syslog host had been set prior to the panic. This is especially useful if multiple panics occur, so that each is recorded.

B.14.2.3 Deleting the Address of the System Log Host

This command allows the user to delete the address of the host to which the switch's system messages are being logged. Enter the following parameters:

> **NOTE:** This command is only available on the local switch.

```
localhost::configuration system syslog> delete

Remove 192.88.243.118 as remote syslog host [n]? y
Remote syslog host removed. Syslog messages will not be sent.
Syslog Facility: local7
```

The switch prompts the user to confirm that the address should be deleted. Entering `y` causes the switch to delete the address, as shown above. If you do not want the address to be deleted, enter `n` or press `<RETURN>` and you will be sent back to the `syslog` prompt.

Once the host's address has been deleted, the switch's system messages are no longer logged until a new host address is set. However, if a facility has been assigned to the switch, that facility assignment remains intact. So if you assign another host without changing the facility, the new host will list this switch's syslog messages under the same facility. For example,

```
localhost::configuration system syslog> show
Remote Syslog Host: 169.144.48.41
Syslog Facility: local7

localhost::configuration system syslog> delete
Remove 169.144.48.41 as remote syslog host [n]? y
Remote syslog host removed. Syslog messages will not be sent.

localhost::configuration system syslog> show
No remote syslog host set. Syslog messages will not be sent.
Syslog Facility: local7

localhost::configuration system syslog> set 204.95.89.84
Remote Syslog Host: 204.95.89.84

localhost::configuration system syslog> show
Remote Syslog Host: 204.95.89.84
Syslog Facility: local7
```

B.14.2.4 Turning Off or Turning On System Log Messages to the Console

On a 9A000, SFCS-200WG, an SFCS-200BX, and an SFCS-1000, the system log messages may be directed to three places: to syslog, to the console, and to a remote host. These types of switches are defaulted to send log messages to both the console and to syslog. The `console` command can be used in different ways.

To display whether or not log messages are being output to the console, enter `console` without any arguments as follows:

```
localhost::configuration system syslog> console
```

If the messages are being sent to the console, the following message is shown:

```
Syslog console output is currently ON.
```

If the messages are not being directed to the console, the following message is displayed:

```
Syslog console output is currently OFF.
```

To enable log messages to be sent to the console, or to stop log messages from being output to the console, enter `console` with the appropriate argument as follows:

```
localhost::configuration system syslog> console [enable | disable]
```

enable Indicates that all log messages will be written to the console.

disable Indicates that no log messages will be written to the console.

B.14.3 AMI Timeout Configuration Command

This command lets the user set the amount of time of non-activity after which an AMI session will time out. The user can get to this level by entering timeout at the system level. Enter the following parameters:

```
localhost::configuration system> timeout [<minutes>]
```

minutes Indicates the number of minutes of non-activity after which an AMI session will time out and exit the user out of the session. The default is 5 minutes. To configure the switch so that an AMI session does not time out, enter 0. It is displayed as off when you use configuration system show.

B.14.4 Configuring the Units for UPC Contracts

This command allows the user to change the type of units that are being used when configuring and displaying UPC contracts. The user can get to this level by entering units at the system level. Enter the following parameters:

```
localhost::configuration system> units (cps | kbps)
```

units Using cps indicates that UPC contracts are being configured and displayed in cells per second. Using kbps means that UPC contracts are being configured and displayed in kilobits per second. The default is cps.

B.15 Topology Configuration Commands

These commands allow the user to manage the *ForeThought* PNNI topology information and the SPANS topology information of the switch fabric. The user can display the list of available subcommands by typing `topology ?` at the `configuration` level.

```
localhost::configuration> topology ?
  forepnni>        spans>
```

B.15.1 *ForeThought* PNNI Configuration Commands

These commands allow the user to modify various aspects of *ForeThought* PNNI on a switch. The user can display the list of available subcommands by typing `forepnni ?` at the `topology` level.

```
localhost::configuration topology> forepnni ?
  prefix           border           swmask           pgmask
  hello            nsapindication   staticupdate     maxhop
  propmult         minthresh        vcmark           show
```

B.15.1.1 Setting the *ForeThought* PNNI Switch Prefix

When using *ForeThought* PNNI, a switch fabric is identified by a variable length NSAP switch prefix which ranges in length from 0 to 13 bytes. This command allows the user to set the *ForeThought* PNNI prefix on the switch. Enter the following parameters:

```
localhost::configuration topology forepnni> prefix <prefix>
```

prefix Indicates the default NSAP prefix for this ATM switch that is used in the ILMI address registration message and in the hello indication SPANS-NNI message.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.1.2 Changing the *ForeThought* PNNI Border Switch Functionality

A switch that has a link to another switch that belonging to a different peergroup is considered a border switch. A border switch advertises reachability to its peergroup to switches outside of its peergroup, but it does not share its peergroup's topology with the other switches. You should enable border switch functionality on all switches that are on the outside edges of all of the peergroups that you have established. This command lets the user designate whether or not this switch will act as a *ForeThought* PNNI border switch. Enter the following parameters:

```
localhost::configuration topology forepnni> border(enable | disable)
```

enable | disable Entering `enable` (and rebooting) means that this switch will act as a *ForeThought* PNNI border switch. Entering `disable` (and rebooting) means that this switch will not act as a *ForeThought* PNNI border switch.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.1.3 Setting the *ForeThought* PNNI Switch Prefix Mask

This command allows the user to select the *ForeThought* PNNI switch prefix mask value. Enter the following parameters:

```
localhost::configuration topology forepnni> swmask <mask>
```

mask Indicates the mask that gives the number of leading bits in the switch prefix used to aggregate the addresses that belong to the switch in *ForeThought* PNNI. The default switch prefix mask value is 104.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.1.4 Setting the *ForeThought* PNNI Peergroup Mask

A peergroup mask is the length (in the number of bits) of the peergroup ID of a switch. This command enables the user to set the *ForeThought* PNNI peergroup mask value. This value should be the same for all members of a peergroup. Enter the following parameters:

```
localhost::configuration topology forepnni> pgmask <mask>
```

mask Indicates the mask that gives the number of leading bits in the switch prefix used to aggregate the addresses that belong to this *ForeThought* PNNI peergroup. The default peergroup mask value is 0.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.1.5 Setting the Hello Indication Interval

Hello indication messages are the "keep alive" messages that two switches send to one another to verify their existence. This command lets the user change the interval for *ForeThought* PNNI hello indication messages. Enter the following parameters:

```
localhost::configuration topology forepnni> hello <msec>
```

msec Indicates the period of time, in milliseconds, between transmissions of hello indication messages. The default value is 500 milliseconds.

B.15.1.6 Setting the NSAP Indication Interval

NSAP indication messages are those messages that update topology information between any two switches. This command allows the user to select the interval for *ForeThought* PNNI NSAP indication messages. Enter the following parameters:

```
localhost::configuration topology forepnni> nsapindication <msec>
```

msec Indicates the period of time, in milliseconds, between transmissions of NSAP indication messages. The default value is 10,000 milliseconds.

B.15.1.7 Setting the Static Route Update Indication Interval

Static route update indication messages are refresh messages that update topology information about static routes. This command enables the user to set the interval for *ForeThought* PNNI static route indication messages. Enter the following parameters:

```
localhost::configuration topology forepnni> staticupdate <msec>
```

msec Indicates the period of time, in milliseconds, between transmissions of static route update indication messages. The default value is 10,000 milliseconds.

B.15.1.8 Setting the Maximum Hop Count

By setting a maximum hop count, the user tells the switch to consider only those paths that have less than or equal to the number of hops specified when setting up a connection. If a connection is routed using a path with a large hop count, there is a greater chance that the connection may experience congestion and be delayed or discarded. This command lets the user set the maximum hop count for the NSAP router. Enter the following parameters:

```
localhost::configuration topology forepnni> maxhop <hops>
```

hops Indicates the maximum number of hops to use when routing a connection for the NSAP router. The default value is 20 hops.

B.15.1.9 Setting the Proportional Multiplier

This command enables the user to set the proportional multiplier for the NSAP router. The proportional multiplier is expressed as a percentage of Available Cell Rate (ACR) on any given link in the network. If the change in percentage of the ACR on any given link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant. The topology tables are then updated accordingly for that link. Enter the following parameters:

```
localhost::configuration topology forepnni> propmult <percentage>
```

percentage Indicates the threshold, entered as a percentage, above which you consider the change in ACR on any link to be significant. The default value is 20%.

NOTE: If you modify this value, you should modify it on all switches in the network.

B.15.1.10 Setting a Minimum Threshold for NSAP Updates

The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which product of the ACR and the proportional multiplier is a very small number. The minimum threshold is used to prevent excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. Enter the following parameters:

```
localhost::configuration topology forepnni> minthresh <minthresh>
```

minthresh Indicates the minimum threshold bandwidth value for triggering NSAP updates, entered in kilobits per second. The default value is 50 kilobits per second.

B.15.1.11 Setting a Minimum Virtual Channel Mark

When the number of available virtual channels on a path drops to zero, a link state update is sent out to advertise that there are no more VCs available for use on this path. When the number of VCs indicated by the vcmark is available for use on this path again, another link state update is sent out to advertise that there are VCs available for use on this path once again. This command lets the user set the vcmark, which is the minimum number of virtual channels that need to be to available on a path to make that path usable again. Enter the following parameters:

```
localhost::configuration topology forepnni> vcmark <vcmark>
```

vcmark Indicates the minimum number of virtual channels that need to be to available on a path to make that path usable. The default value is 20 VCs.

B.15.1.12 Displaying *ForeThought* PNNI Parameters

This command let the user display all of the *ForeThought* PNNI topology parameters. Enter the following parameters:

```
localhost::configuration topology forepnni> show

Switch NSAP prefix                    0x47.0005.80.ffe100.0000.f215.0df6
Switch Prefix Mask                     104
Peer Group Mask                        0

Hello Indication Interval              500 msec
NSAP Indication Interval               10000 msec
Static Route Update Interval           10000 msec
Max hop count for NSAP router          20 hops
Proportional Multiplier                20 %
Minumum Threshold for NSAP updates     50 Kbps
Minimum VC level                       20

FORE PNNI border switch functionality is disabled
```

The fields in this display are defined as follows:

Switch NSAP prefix	Displays the switch's NSAP prefix.
Switch Prefix Mask	Shows the switch prefix mask value of high-order bits to use for aggregating addresses on the switch for routing purposes.
Peer Group Mask	Lists the peergroup mask value of high-order bits to use for aggregating addresses on the switch for routing purposes.
Hello Indication Interval	Displays the period of time between transmissions of hello indication messages, in milliseconds.
NSAP Indication Interval	Shows the period of time between transmissions of NSAP indication messages, in milliseconds.
Static Route Update Interval	Lists the period of time between transmissions of static route update indication messages, in milliseconds.
Max hop count for NSAP router	Displays the maximum number of hops to use when routing a connection for the NSAP router.
Proportional Multiplier	Shows the threshold, in percentage, above which the change in ACR on any link is considered to be significant.
Minimum Threshold for NSAP updates	Lists the minimum threshold bandwidth value for triggering NSAP updates, in kilobits per second.
Minimum VC level	Lists the minimum number of VCs that need to be available on a path to make that path usable again after the number of available VCs has dropped to zero.
FORE PNNI border switch functionality is disabled	If this functionality is `enabled`, this switch acts as a *ForeThought* PNNI border switch. If this functionality is `disabled`, this switch does not act as a *ForeThought* PNNI border switch.

B.15.2 SPANS Topology Configuration Commands

These commands allow the user to modify various aspects of SPANS-NNI on a switch. The user can display the list of available subcommands by typing spans ? at the topology level.

```
localhost::configuration topology> spans ?
  border           area             show
```

NOTE: In order for part of a FORE ATM cloud to be a hierarchical SPANS area, all switches in that cloud must be running software version 4.0. Otherwise, SPANS connectivity between the hierarchical area and the SPANS area will be lost.

B.15.2.1 Setting the SPANS-NNI Border Switch Functionality

A switch that has a link to another switch belonging to a different SPANS area is considered a border switch. A border switch advertises reachability to its area to switches outside of its area, but it does not share its area's topology with the other switches. You should enable border switch functionality on all switches that are on the outside edges of all of the areas that you have established. This command lets the user designate whether or not this switch will act as a SPANS-NNI border switch. Enter the following parameters:

```
localhost::configuration topology spans> border (enable | disable)
```

enable | disable Entering enable (and rebooting) means that this switch will be a SPANS border switch. Entering disable (and rebooting) means that this switch will not be a SPANS border switch.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.2.2 Setting the SPANS Area ID

This command enables the user to set the SPANS area ID. Enter the following parameters:

```
localhost::configuration topology spans> area <area>
```

area Indicates the ID of the area in the SPANS routing hierarchy to which this switch belongs. This number goes into the most significant byte of the SPANS ATM address. The default area ID for all switches is 242 in decimal.

NOTE: The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

B.15.2.3 Displaying SPANS-NNI Parameters

This command let the user display all of the SPANS-NNI topology parameters. Enter the following parameters:

```
localhost::configuration topology spans> show
SPANS Area ID           242
SPANS NNI border switch functionality is disabled
```

The field in this display is defined as follows:

SPANS Area ID Shows the ID of the area in the SPANS routing hierarchy to which this switch belongs. This number goes into the most significant byte of the SPANS ATM address.

SPANS NNI border switch functionality is disabled If this functionality is enabled, this switch is a SPANS-NNI border switch. If it is disabled, this switch is not a SPANS-NNI border switch.

B.16 UNI 3.0 Configuration Commands

These commands let the user configure UNI 3.0 signalling paths. The user can display the list of available subcommands by typing uni30 ? at the configuration level.

```
localhost::configuration> uni30 ?
  delete       new            show
```

B.16.1 Deleting a UNI 3.0 Signalling Path

This command allows the user to delete an existing UNI 3.0 signalling path. Enter the following parameters:

```
localhost::configuration uni30> delete <port> <vpi>
```

port Indicates the port number on which the UNI 3.0 signalling path is to be deleted.

vpi Indicates the number of the UNI 3.0 signalling path to be removed.

B.16.2 Creating a UNI 3.0 Signalling Path

This command enables the user to create a UNI 3.0 signalling path.

NOTE: Before a UNI 3.0 signalling path can be created on a given VPI, an originating and a terminating path must exist for that same VPI.

Enter the following parameters:

```
localhost::configuration uni30> new <port> <vpi> [-ilmi (up | down)] [(user | network)]
    [(publicUNI | auto | IISP)]

advanced options:
     [-sigvci <vci>] [-ilmivci <vci>]
     [-minvci <vci>] [-maxvci <vci>]
     [-sigbw <Kbps>] [-ilmibw <Kbps>]
```

port Indicates the port number on which the UNI 3.0 signalling path is to be created.

vpi Indicates the number of the UNI 3.0 signalling path to be created.

-ilmi (up | down) Enables ILMI NSAP registration for this port (only when a host is connected). The default is `up`. However, this parameter must be set to `down` if you set the last parameter to IISP.

user | network Indicates the switch user side or network side. If the connection is to a host, choose `network`. If the connection is to another switch, one switch must be `user` and the other switch must be `network`.

publicUNI | auto | IISP Indicates the configuration type. `PublicUNI` means that this link is used between this switch and a public switch. `Auto` means that the operation type is determined dynamically. IISP is used for switch-to-switch signalling for static NNI routes. The default is `auto`.

The advanced options are as follows:

-sigvci <vci> Indicates the VCI to use for UNI 3.0 signalling messages. The default reserved VCI is `5`.

-ilmivci <vci> Indicates the VCI to use for ILMI signalling messages. The default reserved VCI is 16.

-minvci <vci> Indicates the bottom number for the range of VCIs to be reserved for UNI 3.0 signalling messages. The default is 32.

-maxvci <vci> Indicates the top number for the range of VCIs to be reserved for UNI 3.0 signalling messages. The default is 511.

-sigbw <Kbps> Indicates the amount of bandwidth that will be reserved on the VCI for UNI 3.0 signalling messages (VCI 5). The default is 0 Kbps.

-ilmibw <Kbps> Indicates the amount of bandwidth that will be reserved on the VCI for ILMI messages (VCI 16). The default is 0 Kbps.

B.16.3 Displaying UNI 3.0 Signalling Paths

This command allows the user to list an individual switch fabric's current UNI 3.0 signalling path information. Enter the following parameters:

```
localhost::configuration uni30> show
 Port  VPI  State ILMI   UNIside  Type  OperType   RemoteAddress
 1C1     0  down  down   network  auto  privateUNI
 1C2     0  down  down   network  auto  privateUNI
 1C3     0  down  down   network  auto  privateUNI
 1C4     0  down  up     network  auto  privateUNI 169.144.64.108
 1D2     0  down  down   network  auto  privateUNI
 1D3     0  down  down   network  auto  privateUNI
 1D4     0  down  down   network  auto  privateUNI
 1CTL    0  up    down   network  auto  privateUNI
```

The fields in this display have the following meaning:

Port Indicates the port number of the network modules that are currently installed in the switch. The 1 indicates that it is the first switch fabric. The letter indicates the position of the network module in the switch. The 1, 2, 3, 4 indicate the specific port number. CTL indicates the control port which is a logical (not physical) location where cells that are directed to the switch itself are sent.

VPI Shows the UNI 3.0 signalling path.

State Shows the current state of the port. If the state is `up`, then this port is functional. This is the normal state for a port that is connected to another Cabletron Systems' ATM switch or host. If the state is `down`, then this port is not functional. This can be due to a lack of a physical connection or due to a software problem.

ILMI `Up` means that ILMI NSAP registration is enabled for this port (only when a host is connected). `Down` means that ILMI NSAP registration is not enabled for this port.

UNIside Indicates the switch user side or network side. If the connection is to a host, `network` is displayed. If the connection is to another switch, one switch must be `user` and the other switch must be `network`.

Type Displays the type of configuration for this port. Can be publicUNI, auto, or IISP. The default is auto.

OperType Shows the signalling type being used on this port.

RemoteAddress Shows the IP address of the remote endstation, if it is available.

To show advanced UNI 3.0 signalling path information about all of the ports on an individual switch fabric, enter the following parameters:

```
localhost::configuration uni30> show advanced
 Port  VPI SigVCI ILMIVCI SigAAL MinVCI MaxVCI SigBW ILMIBW
 1C1     0      5      16      5     32    511     0      0
 1C2     0      5      16      5     32    511     0      0
 1C3     0      5      16      5     32    511     0      0
 1C4     0      5      16      5     32    511     0      0
 1D2     0      5      16      5     32    511     0      0
 1D3     0      5      16      5     32    511     0      0
 1D4     0      5      16      5     32    511     0      0
 1CTL    0      5      16      5     32    511     0      0
```

The fields in this display have the following meanings:

Port Indicates the port number of the network modules that are currently installed in the switch. The 1 indicates that it is the first switch fabric. The letter indicates the position of the network module in the switch. The 1, 2, 3, 4 indicate the specific port number. CTL indicates the control port which is a logical (not physical) location where cells that are directed to the switch itself are sent.

VPI Shows the UNI 3.0 signalling path.

SigVCI Indicates the VCI on the UNI 3.0 signalling path. The default reserved VCI is 5.

IMLIVCI Lists the VCI to be used for ILMI signalling messages. The default reserved VCI is 16.

SigAAL Shows the AAL type being used for this connection. The default is AAL5.

MinVCI Displays the bottom number for the range of VCIs to be reserved for UNI 3.0 signalling messages. The default is 32.

MaxVCI Displays the top number for the range of VCIs to be reserved for UNI 3.0 signalling messages. The default is 511.

SigBW Shows the amount of bandwidth that is reserved on the UNI 3.0 signalling VCI. The default is 0.

ILMIBW Lists the amount of bandwidth that is reserved on the ILMI VCI. The default is 0.

To list UNI 3.0 signalling path information for a specific port, (for example, port 1C1), enter the following parameters:

```
localhost::configuration uni30> show 1c1
 Port  VPI  State ILMI   UNIside  Type  OperType   RemoteAddress
 1C1     0  down  down   network  auto  privateUNI
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying UNI 3.0 signalling path information on all of the ports on an individual switch fabric.

To list UNI 3.0 signalling path information for a specific port and path, (for example, port 1C1 and VPI 0), enter the following parameters:

```
localhost::configuration uni30> show 1c1 0
 Port  VPI  State ILMI   UNIside  Type  OperType   RemoteAddress
 1C1     0  down  down   network  auto  privateUNI
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying UNI 3.0 signalling path information on all of the ports on an individual switch fabric.

To list advanced UNI 3.0 signalling path information for a specific port and path, (for example, port 1C1 and VPI 0), enter the following parameters:

```
localhost::configuration uni30> show 1c1 0 advanced
 Port  VPI SigVCI ILMIVCI SigAAL MinVCI MaxVCI SigBW ILMIBW
 1C1     0      5      16      5     32    511     0      0
```

The fields in this display are defined in the same manner as those listed previously in the example for displaying advanced UNI 3.0 signalling path information on all of the ports on an individual switch fabric.

B.17 Usage Parameter Control Configuration Commands

These commands let the user create or delete usage parameter control (UPC) traffic contracts. The user can display the list of available subcommands by typing `upc ?` at the `configuration` level.

```
localhost::configuration> upc ?

delete          new             show
```

B.17.1 Deleting a UPC Traffic Contract

This command allows the user to delete an existing UPC contract. Enter the following parameters:

```
localhost::configuration upc> delete <index>
```

index Indicates the integer index that refers to a specific traffic contract.

B.17.2 Creating a UPC Traffic Contract

This command allows the user to create a UPC contract. Enter the following parameters:

```
localhost::configuration upc> new <index> [<UPC>] [-cdvt <us>] [aal5epd] [-name <name>]
Where UPC is one of the following combinations of traffic parameters:
    ubr
    cbr <pcr01>
    cbr0 <pcr0> <pcr01> [tag]
    vbr <pcr01> <scr01> <mbs01>
    vbr0 <pcr01> <scr0> <mbs0> [tag]
```

index Enter an integer index that will be used to refer to this specific traffic contract.

UPC Indicates one of the types of traffic contracts shown above. The parameters in these contracts are defined as follows:

ubr Indicates UBR traffic.

cbr Indicates CBR traffic.

cbr0 Indicates CBR0 traffic.

vbr Indicates VBR traffic.

vbr0 Indicates VBR0 traffic.

<pcr0> Indicates the peak cell rate for cells with CLP = 0.

<pcr01> Indicates the peak cell rate for all cells.

<scr0> Indicates the sustainable cell rate for cells with CLP = 0.

<scr01> Indicates the sustainable cell rate for all cells.

<mbs0> Indicates the maximum burst size for cells with CLP = 0.

<mbs01> Indicates the maximum burst size for all cells.

NOTE: The units for pcr0, pcr01, scr0, scr01, mbs0, and mbs01 are specified either in cells per second or in kilobits per second, depending on what you have configured for configuration system units. To display the current setting, use configuration system show. The default is cps (cells per second).

[tag] Entering `tag` means that non-conforming CLP = 0 cells are tagged. Otherwise, they are dropped. `Drop` is the default. This option only applies to the PCR0 parameter of the CBR0 contract and to the SCR0 and MBS0 parameters of the VBR0 contract.

-cdvt <us> Indicates the Cell Delay Variation Tolerance (CDVT) associated with the peak cell rates, specified in microseconds.

aal5epd Entering `aal5epd` means that Early Packet Discard is enabled on this connection. If `aal5epd` is not entered, then Early Packet Discard is not enabled on this connection.

-name <name> Indicates the user-defined name associated with this UPC traffic contract. This helps the user remember for what traffic type this specific contract is used. If the user does not specify a name, a default name that relates to this type of traffic contract is assigned automatically.

The following is an example of how to create a UPC contract:

```
localhost::configuration upc> new 5 vbr0 500 200 250 -cdvt 1000 aal5epd -name vbr0_upc
```

This example specifies a contract named "vbr0_upc", which is a VBR0 contract with an index of 5, a pcr01 of 500 cells/sec (or kbps), an scr0 of 200 cells/sec (or kbps), an mbs0 of 250 cells (or kilobits), a CDVT of 1,000 microseconds, and EPD enabled.

NOTE: For more information regarding traffic contracts, please refer to Table 5-7 in the ATM Forum UNI 3.0 Specification.

B.17.3 Displaying the UPC Traffic Contracts

This command lets the user display all of the UPC contracts. Enter the following parameters:

```
localhost::configuration upc> show
Index  PCR01  SCR01  MBS01   PCR0   SCR0   MBS0   CDVT Act  EPD Name
0                                                   drop no  default_ubr
1         50                  25                    tag  no  cbr0_upc
2        500                       200    250   1000 drop yes vbr0_upc
```

Index Shows the UPC contracts listed by index number.

PCR01 Shows the peak cell rate for all cells for this contract.

SCR01 Lists the sustainable cell rate for all cells in this contract.

MBS01 Lists the maximum burst size for all cells in this contract.

PCR0 Shows the peak cell rate for cells with CLP = 0 for this contract.

SCR0 Shows the sustainable cell rate for cells with CLP = 0 for this contract.

MBS0 Shows the maximum burst size for cells with CLP = 0 for this contract.

CDVT Shows the Cell Delay Variation Tolerance (CDVT) associated with the peak cell rates, in microseconds.

Act `Tag` means that non-compliant CLP = 0 cells are tagged. `Drop` means that non-compliant cells are dropped. This option only applies to cells measured by the PCR0 parameter of the CBR0 contract and to cells measured by the SCR0 and MBS0 parameters of the VBR0 contract.

EPD `Yes` means that EPD is enabled on this connection. `No` means that EPD is not enabled on this connection.

Name Shows the user-defined name associated with this UPC traffic contract.

NOTE: The units for `PCR0`, `PCR01`, `SCR0`, `SCR01`, `MBS0`, and `MBS01` are shown either in cps or in kbps, depending on what you have configured for `configuration system units`. To display the current setting, use `configuration system show`. The default is `cps`.

B.18 Virtual Channel Configuration Commands

These commands let the user configure permanent virtual channels (PVCs). The user can display the list of available subcommands by typing `vcc ?` at the `configuration` level.

```
localhost::configuration> vcc ?

delete          modify          new             show
```

B.18.1 Deleting a Virtual Channel

This command allows the user to delete an existing permanent virtual channel. Enter the following parameters:

```
localhost::configuration vcc> delete <iport> <ivpi> <ivci> <oport> <ovpi> <ovci>
```

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

ivci Indicates the incoming virtual channel number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

ovci Indicates the outgoing virtual channel number.

B.18.2 Modifying a Virtual Channel

This command enables the user to modify the UPC traffic contract of an existing permanent virtual channel without having to delete the channel and then recreate it. Enter the following parameters:

```
localhost::configuration vpc> modify <port> <vpi> <vci> -upc <index>
```

port Indicates the port number.

vpi Indicates the virtual path number.

vci Indicates the virtual channel number.

-upc<index> Indicates the integer index that refers to a specific UPC traffic contract. If no index is specified, then no traffic policing will take place on this VCI. It is assigned a UPC index of 0, and all traffic on this VCI is treated as UBR traffic. This is the default.

B.18.3 Creating a Virtual Channel

This command enables the user to add a new permanent virtual channel. Enter the following parameters:

```
localhost::configuration vcc> new <iport> <ivpi> <ivci> <oport> <ovpi> <ovci> [-upc <index>]
```

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

ivci Indicates the incoming virtual channel number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

ovci Indicates the outgoing virtual channel number.

-upc<index> Indicates the integer index that refers to a specific UPC traffic contract. If no index is specified, then no traffic policing will take place on this VCI. It is assigned a UPC index of 0, and all traffic on this VCI is treated as UBR traffic. This is the default.

The following is an example of how to create a virtual channel on an SFCS-1000. To create a vcc going in port 2A1, vpi 0, vci 100 on the switch board installed in slot 2 and going out port 4B1, vpi 0, vci 100 on the switch board installed in slot 4, enter the following:

```
localhost::configuration vcc> new 2a1 0 100 2e4 0 100
localhost::configuration vcc> new 2e4 0 100 2a1 0 100

localhost::configuration vcc> new 4b1 0 100 4e2 0 100
localhost::configuration vcc> new 4e2 0 100 4b1 0 100
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

B.18.4 Displaying the Virtual Channel Configuration

This command allows the user to display existing virtual channels. The user can display either all of the existing virtual channels on an individual switch fabric or all of the existing virtual channels on a specific port. To list all of the existing permanent virtual channels on an individual switch fabric, enter the following parameters:

```
localhost::configuration vcc> show
 Input           Output
 Port  VPI  VCI  Port  VPI  VCI  UPC  Protocol
 1C1     0    5  1CTL    0   34       uni30
 1C1     0   14  1CTL    0   33    0  spans
 1C1     0   15  1CTL    0   32       spans
 1C1     0   16  1CTL    0   58       uni30
 1C2     0    5  1CTL    0   37       uni30
 1C2     0   14  1CTL    0   36    0  spans
 1C2     0   15  1CTL    0   35       spans
 1C2     0   16  1CTL    0   59       uni30
 1C3     0    5  1CTL    0   40       uni30
 1C3     0   14  1CTL    0   39    0  spans
 1C3     0   15  1CTL    0   38       spans
 1C3     0   16  1CTL    0   60       uni30
 1C4     0    5  1CTL    0   43       uni30
 1C4     0   14  1CTL    0   42    0  spans
 1C4     0   15  1CTL    0   41       spans

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

Input Port Lists the incoming port number of the virtual channel.

Input VPI Shows the incoming virtual path number.

Input VCI Indicates the incoming virtual channel number.

Output Port Lists the outgoing port number of the virtual channel.

Output VPI Shows the outgoing virtual path number.

Output VCI Indicates the outgoing virtual channel number.

UPC Shows the integer index that refers to the specific UPC traffic contract assigned to this VCI.

Protocol Displays what type of protocol is running on this channel, which can be `spans`, `pvc`, `uni30` or `spvc`.

To list all of the existing permanent virtual channels on a specific port, (for example, port 1C1), enter the following parameters:

```
localhost::configuration vcc> show 1c1
 Input           Output
 Port  VPI  VCI  Port  VPI  VCI  UPC  Protocol
 1C1     0    5  1CTL    0   34       uni30
 1C1     0   14  1CTL    0   33    0  spans
 1C1     0   15  1CTL    0   32       spans
 1C1     0   16  1CTL    0   58       uni30
```

The fields in this display are defined in the same manner as those listed above in the example for all of the permanent virtual channels on an individual switch fabric.

To list all of the existing permanent virtual channels on a specific port and path, (for example, port 1C1 and VPI 0), enter the following parameters:

```
localhost::configuration vcc> show 1c1 0
 Input           Output
 Port  VPI  VCI  Port  VPI  VCI  UPC  Protocol
 1C1     0    5  1CTL    0   34       uni30
 1C1     0   14  1CTL    0   33    0  spans
 1C1     0   15  1CTL    0   32       spans
 1C1     0   16  1CTL    0   58       uni30
```

The fields in this display are defined in the same manner as those listed above in the example for all of the permanent virtual channels on an individual switch fabric.

To list all of the existing permanent virtual channels on a specific port, path, and channel (for example, port 1C1, VPI 0, and VCI 16), enter the following parameters:

```
localhost::configuration vcc> show 1c1 0 16
 Input           Output
 Port  VPI  VCI  Port  VPI  VCI  UPC  Protocol
 1C1     0   16  1CTL    0   58       uni30
```

The fields in this display are defined in the same manner as those listed above in the example for all of the permanent virtual channels on an individual switch fabric.

B.19 Virtual Path Configuration Commands

These commands let the user configure virtual paths. The user can display the list of available subcommands by typing `vpc ?` at the `configuration` level.

```
localhost::configuration> vpc ?

delete          modify          new             show
```

B.19.1 Deleting a Virtual Path

This command lets the user delete an existing virtual path. Enter the following parameters:

```
localhost::configuration vpc> delete <port> <vpi> (<port> <vpi> | term | orig)
```

port Indicates the number of the input port on which the virtual path is to be deleted.

vpi Indicates the number of the input VPI to be deleted.

port Used to delete a through path. Indicates the number of the output port on which the through path is to be deleted.

vpi Used to delete a through path. Indicates the number of the output VPI to be deleted.

term Used to delete a terminating path.

orig Used to delete an originating path.

NOTE: Virtual Path 0 cannot be deleted on any of the connections to the intra-fabric ports on an SFCS-1000.

NOTE: Before deleting a virtual path, you must first delete all VCCs which use that path.

B.19.2 Modifying a Virtual Path

This command lets the user modify the UPC contract of a through path or modify the amount of reserved bandwidth for an existing virtual path without having to delete the path and then recreate it. Enter the following parameters:

```
localhost::configuration vpc> modify <iport> <ivpi> <oport> <ovpi> -upc <index>
or:
localhost::configuration vpc> modify <port> <vpi> (term | orig) -reserved <Kbs>
```

The following parameters are used to modify the UPC contract on a through path:

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

-upc <index> Indicates the integer index that refers to a specific UPC traffic contract.

The following parameters are used to modify the reserved bandwidth on a terminating or originating path:

port Indicates the port number to be modified.

vpi Indicates the virtual path number to be modified.

term | orig Indicates whether the virtual path to be modified is terminating or originating.

-reserved<Kbs> Indicates the amount of bandwidth, specified in Kbps, to be reserved for this path.

B.19.3 Creating a Virtual Path

This command lets the user add a virtual path in three different ways. The first line is used for creating a through path. The second line is used for creating a terminating path. The third line is used for creating an originating path. Both through paths and originating paths have advanced options which may be used in combination with the required parameters for that type of path. Enter the following parameters:

```
localhost::configuration vpc> new <iport> <ivpi> <oport> <ovpi> [-upc <index>]
advanced options:
[-shapeivpi <vpi>] (through)
```

or

```
localhost::configuration vpc> new <port> <vpi> term [-reserved <Kbs>] [-maxvci <maxvci>]
```

or

```
localhost::configuration vpc> new <port> <vpi> orig [-reserved <Kbs>] [-maxvci <maxvci>]
advanced options:
[-shapeovpi <vpi>] [-vbrob <percent>] [-vbrbuffob <percent>] (orig)
```

The following parameters make up a through path:

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

-upc <index> Indicates the integer index that refers to a specific UPC traffic contract. If no index is specified, then no traffic policing will take place on this VPI. It is assigned a UPC index of 0, and all traffic on this VPI is treated as UBR traffic. This is the default.

The following is an advanced option for through paths:

-shapeivpi <vpi> Indicates the incoming VPI. This can only be set when creating a through path. When the traffic shaping port is not the port connected to the WAN, a through path must be created from the WAN port to the traffic shaping port. Cells arrive from the network at the traffic shaping port with this value equal to the VPI of the terminating path at the traffic shaping port.

The following parameters make up a terminating path:

port Indicates the incoming port number.

vpi Indicates the incoming virtual path number.

term Specifies that the virtual path to be created is a terminating path. If this option is not used, an elastic path is created. Elastic paths allocate and de-allocate bandwidth for their channels from the link.

-reserved <Kbs> Indicates the amount of bandwidth specified in Kbps that the user wants to reserve on this path.

-maxvci <maxvci> Indicates the maximum number of channels that can be created on this path.

The following parameters make up an originating path:

port Indicates the incoming port number.

vpi Indicates the incoming virtual path number.

orig Specifies that the virtual path to be created is an originating path.

-reserved <Kbs> Indicates the amount of bandwidth specified in Kbps that the user wants to reserve on this path.

-maxvci <maxvci> Indicates the maximum number of channels that can be created on this path.

The following are advanced options for originating paths:

-shapeovpi <vpi> Indicates the output port of a traffic shaping originating path. Setting this value configures traffic shaping on the originating path. Cells bound for the network leave the traffic shaping port with this VPI. When the traffic shaping port is the WAN port, this value equals the input VPI of the originating path. If the traffic shaping port is not the WAN port, this value equals the input VPI of the through path from the shaping port to the WAN port.

-vbrob <percent> Indicates the bandwidth overbooking level assigned to this path, specified as a percentage. Enter an integer value from 1 to 500. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

-vbrbuffob <percent> Indicates the buffer overbooking level assigned to this path, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

NOTE: If you want to shape traffic on more than two ports on a given network module, it is recommended that you set the traffic memory model to model number 5 for that network module. Please see the sections, Displaying Traffic Models for a Network Module, and, Setting Traffic Models on a Network Module, found earlier in this chapter for information about how to do this.

NOTE: Terminating and originating paths cannot be created across the intra-fabric ports on an SFCS-1000; only through paths can be created across the intra-fabric ports as shown in the following example.

The following is an example of how to create a virtual path on an SFCS-1000. To create a through path going in port 2A1, vpi 1 on the switch board installed in slot 2 and going out port 4B1, vpi 1 on the switch board installed in slot 4, enter the following:

```
localhost::configuration vpc> new 2a1 1 2e4 1
localhost::configuration vpc> new 2e4 1 2a1 1

localhost::configuration vpc> new 4b1 1 4e2 1
localhost::configuration vpc> new 4e2 1 4b1 1
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

B.19.4 Displaying Virtual Paths

This command lets the user display existing virtual paths. The user can show either all of the existing virtual paths on an individual switch fabric or all of the existing virtual paths on a specific port. Enter the following parameters:

```
localhost::configuration vpc> show
Input      Output
Port  VPI  Port  VPI   MaxBW     BW MaxVCs  VCs  UPC  Prot
1A1     0  terminate     N/A   0.0K    511    4  N/A  pvc
1A2     0  terminate     N/A   0.0K    511    4  N/A  pvc
1A3     0  terminate     N/A   0.0K    511    4  N/A  pvc
1A4     0  terminate     N/A   0.0K    511    4  N/A  pvc
1A5     0  terminate     N/A   0.0K    511    4  N/A  pvc
1A6     0  terminate     N/A   0.0K    511    4  N/A  pvc
1B1     0  terminate     N/A   0.0K    511    4  N/A  pvc
1B2     0  terminate     N/A   0.0K    511    4  N/A  pvc
1B3     0  terminate     N/A   0.0K    511    4  N/A  pvc
1B4     0  terminate     N/A   0.0K    511    4  N/A  pvc
1CTL    0  terminate     N/A   0.0K    511   36  N/A  pvc
originate  1A1     0     N/A   0.0K    511    4  N/A  pvc
originate  1A2     0     N/A   0.0K    511    4  N/A  pvc
originate  1A3     0     N/A   0.0K    511    4  N/A  pvc
originate  1A4     0     N/A   0.0K    511    4  N/A  pvc
originate  1A5     0     N/A   0.0K    511    4  N/A  pvc
originate  1A6     0     N/A   0.0K    511    4  N/A  pvc

Press return for more, q to quit: q
```

The fields in this display are defined as follows:

Input Port — Shows the number of the input port of the virtual path. Lists originate if it is an originating path.

Input VPI — Shows the input virtual path.

Output Port — Shows the number of the output port of the virtual path. Lists terminate if it is a terminating path.

Output VPI — Shows the output virtual path.

MaxBW — Lists the maximum amount of bandwidth (in megabits/second) that is reserved for the virtual channels using this path. A value of N/A indicates that this path is an elastic path. Elastic paths allocate and deallocate bandwidth for their channels from the link.

BW Shows the amount of bandwidth (in megabits/second) that has been reserved for the virtual channels using this path.

MaxVCs Lists the maximum number of virtual channels that may use this originating or terminating path.

VCs Lists the number of virtual channels that are currently using this originating or terminating path.

UPC Shows the integer index that refers to a specific traffic contract assigned to this through path.

Protocol Lists `pvc` for a permanent virtual circuit, `spans` for a switched virtual circuit, or `uni` for a UNI 3.0 path.

To list all of the advanced options about existing virtual paths, enter the following parameters:

```
localhost::configuration vpc> show advanced
 Input      Output
 Port  VPI  Port  VPI  Shape VBROB BuffOB
 1A1     0  terminate    N/A   N/A    N/A
 1A2     0  terminate    N/A   N/A    N/A
 1A3     0  terminate    N/A   N/A    N/A
 1A4     0  terminate    N/A   N/A    N/A
 1A5     0  terminate    N/A   N/A    N/A
 1A6     0  terminate    N/A   N/A    N/A
 1B1     0  terminate    N/A   N/A    N/A
 1B2     0  terminate    N/A   N/A    N/A
 1B3     0  terminate    N/A   N/A    N/A
 1B4     0  terminate    N/A   N/A    N/A
 1CTL    0  terminate    N/A   N/A    N/A
 originate  1A1     0          100    100
 originate  1A2     0          100    100
 originate  1A3     0          100    100
 originate  1A4     0          100    100
 originate  1A5     0          100    100
 originate  1A6     0          100    100

 Press return for more, q to quit:
```

The fields in this display are defined as follows:

Input Port Shows the number of the input port of the virtual path. Lists `originate` if it is an originating path.

Input VPI Indicates the input virtual path.

Output Port Shows the number of the output port of the virtual path. Lists terminate if it is a terminating path.

Output VPI Shows the output virtual path.

Shape Indicates whether or not traffic shaping has been enabled for this originating path.

VBROB Shows the bandwidth overbooking level assigned to this path, specified as a percentage. Valid values are integers from 1 to 500. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

BuffOB Shows the buffer overbooking level assigned to this path, specified as a percentage. Valid values are integers greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

To list all of the existing virtual paths on a specific port, (for example, port 1A1, enter the following parameters:

```
localhost::configuration vpc> show 1A1
 Input      Output
 Port  VPI  Port  VPI   MaxBW     BW MaxVCs  VCs  UPC  Prot
 1A1     0  terminate     N/A   0.0K    511    4  N/A  pvc
 originate  1A1     0     N/A   0.0K    511    4  N/A  pvc
```

The fields in this display are defined in the same manner as those listed above in the example for all of the virtual paths on an individual switch fabric.

To list all of the existing permanent virtual paths on a specific port and path, (for example, port 1A1 and VPI 0) enter the following parameters:

```
localhost::configuration vpc> show 1A1 0
 Input      Output
 Port  VPI  Port  VPI   MaxBW     BW MaxVCs  VCs  UPC  Prot
 1A1     0  terminate     N/A   0.0K    511    4  N/A  pvc
 originate  1A1     0     N/A   0.0K    511    4  N/A  pvc
```

The fields in this display are defined in the same manner as those listed above in the example for all of the permanent virtual paths on an individual switch fabric.

To list all of the advanced options about existing permanent virtual paths on a specific port and path, (for example, port 1A1 and VPI 0) enter the following parameters:

```
localhost::configuration vpc> show 1A1 0 advanced
 Input      Output
 Port  VPI  Port  VPI  Shape VBROB BuffOB
 1A1     0  terminate    N/A   N/A    N/A
 originate  1A1     0          100    100
```

The fields in this display are defined in the same manner as those listed above in the example for the advanced options for all of the permanent virtual paths on an individual switch fabric.

APPENDIX C AMI Operation Commands

This chapter contains a detailed description of the AMI `operation` commands. The main `operation` menu can be found at the root level. There are several commands available under `operation`. Commands that are submenus are immediately followed by a ">" symbol. Typing `operation ?` at the prompt at the root level displays the `operation` commands for the i960-based switches as follows:

```
localhost::operation> ?
  cdb>             environment>     panic>           date
  flash>           password         upgrade          version
  reboot
```

C.1 Configuration Database (CDB) Operation Commands

These commands allow the user to manage the configuration database (CDB). Typing `cdb ?` at the prompt at the `operation` level displays the `cdb` commands as follows:

```
localhost::operation> cdb ?
backup           init             reset            restore
```

C.1.1 Backing Up the Database

This command lets the user make a backup of the configuration database. On a 9A000, SFCS-200WG, SFCS-200BX, and an SFCS-1000, the remote host to which the file will be backed up must be running the tftp server code.

Since the SCP on a 9A000, SFCS-200WG, SFCS-200BX, and an SFCS-1000 uses tftp to perform the CDB backup, you must first create an empty file in the `/tftpboot` directory on the remote host to receive the CDB. Use the `touch` command to do this. Then, use the `chmod` command to change the permissions of that file so that it will let the switch write the backup CDB to that file.

Perform the following steps to back up your CDB:

1. First, telnet to your remote host and log in.
2. Enter the following commands in sequence:

   cd /tftpboot

   touch <backup file name>

   chmod 777 <backup file name>

3. Then exit from the telnet session.
4. Telnet to the switch and log into AMI.
5. Enter the following command:

   oper cdb backup <host>:/tftpboot/<backup file name>

You should receive the following message:

```
CDB backup was successful
```

Your backup file now resides in the file and on the host you specified.

C.1.2 Initializing the Database

This command lets the user initialize the CDB. The switch asks you to verify this action before it re-initializes the CDB.

Enter the following parameters:

```
localhost::operation cdb> init
This command will re-initialize the CDB and reboot the switch
Do you really want to remove ALL permanent information from
the database INCLUDING the configuration of all the network
interfaces? [n] n

localhost::operation cdb>
```

C.1.3 Resetting the Database

This command enables the user to reset the configuration database. The only information that will be retained is the IP configuration which includes the switch name and interface descriptions. The switch cautions the user that all ATM information will be deleted. The switch then asks the user to confirm that resetting the cdb is the desired action. Enter the following parameters:

```
localhost::operation cdb> reset

********** W A R N I N G **********

This operation resets the switch configuration database.
As a result, the switch control software will be restarted.
You will lose connectivity with the switch while this
operation is progressing.

Are you sure you want to reset the CDB [n]? n
localhost::operation cdb>
```

If the user enters yes to the reset question, the switch responds as follows:

```
Are you sure you want to reset the CDB [n]? y
The switch will restart momentarily.
```

At this point, the switch resets the CDB, closes the user out of all active sessions, and restarts the switch. The user must then log in to AMI again to perform any more actions on the switch.

C.1.4 Restoring the Database

This command allows the user to restore the configuration database. Enter the following parameters:

```
localhost::operation cdb> restore [<host>:]<full path to backup file>
```

host — Specifies the IP address of the host on which the CDB file that is to be restored resides.

full path to backup file — Indicates the full path name of the CDB file that is to be restored.

C.2 Environment Commands

These commands allow the user to monitor the switch's environmental parameters. Typing `environment ?` at the prompt at the `operation` level displays the `cdb` commands as follows:

```
localhost::operation> environment ?
cpu             fabric>          fans             power
temperature
```

C.2.1 CPU Operation

This command lets the user display information about the CPU on a 9A000, SFCS-200WG, an SFCS-200BX, and an SFCS-1000. Enter the following parameters:

```
localhost::operation environment> cpu
 CPU        Type       State
 1X         i960       normal
```

The fields in this display are defined as follows:

CPU Indicates the CPU number.

Type Indicates the type of processor.

State Shows whether or not the current condition of the CPU is normal.

C.2.2 Switch Fabric Operation

These commands allow the user to monitor the temperature of the individual switch fabrics on an SFCS-1000 only. Typing `fabric ?` at the prompt at the `environment` level displays the `cdb` commands as follows:

```
localhost::operation environment> fabric ?
  show              temperature
```

C.2.3 Showing Switch Fabric Temperature Information

This command displays the current temperatures in degrees Celsius of each installed switch fabric on an SFCS-1000, the current state of the temperature sensor, and the current thresholds at which a temperature alarm trips and then later resets. The current temperature and state values are displayed for all installed fabrics, even if a fabric has been shut down by the Common Equipment Card (CEC) because of an over temperature condition. Enter the following parameters:

```
localhost::operation environment fabric> show

Fabric     Deg C      State
1          31         normal
2          28         normal
3          37         normal
4          35         normal
Alarm/trap reset threshold (this fabric): 60 degrees C or lower
Alarm/trap trip threshold (this fabric): 65 degrees C or greater
```

The fields in this display are defined as follows:

Fabric Shows the number of the fabrics currently installed in the switch. Switch fabric 1 is in the slot labeled 1 on the enclosure, switch fabric 2 is in the slot labeled 2 on the enclosure, and so on.

Deg C Lists the current temperature of the switch fabrics in degrees Celsius.

State Displays `overTemp` if an alarm has been tripped because of this sensor, based on the trip and reset values that have been configured. Shows `normal` if otherwise, or if the alarm has reset.

Alarm/trap reset threshold Shows the temperature in ^{o}C at which an overtemperature alarm is reset. For example, if the user sets the reset and trip thresholds to 50 ^{o}C and 60 ^{o}C, respectively, then the alarm will trip at 60 ^{o}C, and it will be reset when the temperature drops back down to 50 ^{o}C.

Alarm/trap trip threshold Shows the temperature in ^{o}C at which an overtemperature alarm is tripped. For example, if the user sets the reset and trip thresholds to 50 ^{o}C and 60 ^{o}C, respectively, then the alarm will trip at 60 ^{o}C, and it will be reset when the temperature drops back down to 50 ^{o}C.

C.2.4 Configuring the Switch Fabric Temperature Thresholds

This command allows the user to set the thresholds at which a temperature alarm will be tripped and then later reset on an SFCS-1000. Any temperature can cause the switch to display a state of `normal` or `overTemp`, depending on the trip and reset thresholds set by the user. For example, a temperature of 55 ^{o}C would show a state of `normal` if the trip threshold was 60 ^{o}C and the switch fabric temperature never reached 60 degrees, but it would show a state of `overTemp` if the switch fabric temperature had reached 60 ^{o}C, and then had dropped to 55 ^{o}C, but had not yet reached a reset threshold set at 50 ^{o}C. Enter the following parameters:

```
localhost::operation environment fabric> temperature <reset threshold> <trip threshold>
```

reset threshold Indicates the temperature in ^{o}C at which an overtemperature alarm is reset. The default is 60 ^{o}C.

trip threshold Indicates the temperature in ^{o}C at which an overtemperature alarm is tripped. The default is 65 ^{o}C.

C.2.5 Fan Operation

This command enables the user to display information about the fans on an SFCS-1000 only. Enter the following parameters:

```
localhost::operation environment> fans
FanBank         FanBankState
1               normal
2               normal
3               normal
4               normal
```

FanBank FanBank corresponds to a single fan, indicating the number of the fan.

FanBankState Displays the current state of the fan. In general, it reads normal. If the fan has malfunctioned, it reads failed.

C.2.6 Power Supply Operation

This command enables the user to display information about power supplies. Enter the following parameters to display information for a DC-powered SFCS-1000:

```
localhost::operation environment> power
 PowerSupply  Type            InputState  OutputState S/N        Version
 1            ps48VDC         normal      normal      107        1
 2            ps48VDC         normal      normal      195        1
```

Enter the following parameters to display information for an SFCS-1000 with tall AC power supplies:

```
localhost::operation environment> power
 PowerSupply  Type            InputState  OutputState
 1            psAutoRangeAC   normal      normal
 2            psAutoRangeAC   normal      normal
```

Enter the following parameters to display information for an SFCS-1000 with short AC power supplies:

```
localhost::operation environment> power
 PowerSupply  Type            InputState  OutputState S/N        Version
 1            psRM1000HA      normal      normal      12         1
 2            psRM1000HA      normal      normal      22         1
```

Enter the following parameters to display information about an SFCS-200BX:

```
localhost::> operation environment> power
 PowerSupply  Type            InputState  OutputState
 1            psAutoRangeAC   normal      normal
 2            psAutoRangeAC   normal      normal
```

Enter the following parameters to display information about an SFCS-200WG:

```
localhost::> operation environment> power
 PowerSupply  Type            InputState  OutputState
 1            psAutoRangeAC   normal      normal
```

The fields in these displays are defined as follows:

PowerSupply On an SFCS-200BX, 1 indicates the left power supply and 2 indicates the right power supply. On an SFCS-1000, 1 indicates the power supply in slot 1 in the chassis and 2 indicates the power supply in slot 2 in the chassis.

Type Displays whether it is an AC or a DC power supply.

InputState Shows if the voltage coming into the power supply is normal or not.

OutputState Shows if the voltage going out of the power supply is normal or not.

S/N Indicates the serial number of the power supply.

Version Lists the power supply's hardware version number.

C.2.7 Temperature Sensor Operation

This command enables the user to display information gathered by the over-temperature sensors. Enter the following parameters:

```
localhost::operation environment> temperature
TemperatureSensor              SensorState
enclosure                      normal
power-supply-A                 normal
power-supply-B                 normal
```

The fields in this display are defined as follows:

TemperatureSensor Indicates where the temperature sensor is located on the unit.

SensorState Shows if the temperature at the specified location is normal or not.

NOTE: An SFCS-200WG only shows an enclosure temperature sensor.

C.2.8 Panic Acknowledgment Commands

On occasion, the SCP may go into a state called panic, in which it reboots, closes a user out of session, or goes into a hung or frozen state as the result of a software bug. When the SCP returns to a normal state and an active session is running again, the first thing the user should do is to use the `operation panic show` AMI command to display information about what happened to the SCP when it panicked. This information helps FORE's Technical Support staff to diagnose the problem. Typing `panic ?` at the prompt at the `operation` level displays the `panic` commands as follows:

```
localhost::operation> panic ?
clear            show
```

C.2.9 Clearing the Panic Flag

This command lets the user clear the panic acknowledgment flag without viewing the contents of the panic dump file. Once the flag is cleared, the user may return to normal operation.

NOTE: Do not clear a panic condition until after you have performed the following three steps.

1. Use the `operation panic show` command to display the contents of the panic file.
2. Cut and paste this panic file information to another file on a host and save that file.
3. Send this information via e-mail to FORE Technical Support along with a description of the events leading up to the panic. Ask the Technical Support staff to open a case for you based on that information. Once you have sent them the information, you may clear the panic.

Enter the following parameters to clear a panic condition:

```
localhost::operation panic> clear
OK.
```

The message above is shown when a panic has been cleared.

The message below is shown when no panic dump file exists (i.e., the SCP did not panic).

```
localhost::operation panic> clear
There is no panic condition to clear.
```

C.2.10 Displaying the Panic Dump File

This command lets the user view the contents of the panic dump file, which contains information about what happened to the SCP when it panicked, without clearing the panic flag. This information can assist FORE's Technical Support staff in troubleshooting the cause of the panic. Once the flag is cleared, the user may return to normal operation.

The following is an example of the kind of message that appears on your console when a panic occurs:

```
****************************************
ForeThought SCP Software
Copyright (C) 1992-1995 FORE Systems, Inc.
All rights reserved.
****************************************
Starting kernel....
PANIC REBOOT! (Complete)
Panicked at Wed Nov 8 17:38:58 1995
with message: Fault
and 10 trace entries.
Use the AMI command `oper panic clear' to clear the panic.
```

At this point, the user must log in to AMI and enter the following parameters to display the panic message:

```
localhost::> oper panic show

The last recorded panic was at Wed Nov  8 17:38:58 1995
Message: [fault]It was for version 3.4.0 revision 1.29 with 2 deltas.
Global registers:
       G0  002C284D   G1  001F0000   G2  00000003   G3  00000000
       G4  002414A0   G5  00000028   G6  000000D8   G7  003368D0
       G8  000001F0   G9  00011D70   G10 00020001   G11 00039B50
       G12 00000002   G13 000000FF   G14 00000000   G15 0000CB60
    Stack trace:
         0. 00245F64
         1. 00039B50
         2. 00245FE4
         3. 00230258
         4. 002414A4
         5. 00229B04
         6. 00229930
         7. 00232E60
         8. 001111A4
         9. 00100088
```

Once the information has been displayed, follow the steps listed in the previous subsection for clearing a panic.

The message below is shown whenever there is no panic dump file.

```
localhost::operation panic> show
There is no panic dump to show you. Thank goodness.
```

NOTE: If the switch panics, the panic file is automatically written to syslog upon reboot, provided that a syslog host had been set prior to the panic. This is especially useful if multiple panics occur, so that each is separately recorded and is not overwritten as they are here. For more information about setting the syslog host, please refer to Appendix B of this manual.

C.3 Displaying and Setting the Date and Time

This command allows the user to display the current date and time on the switch. To display this information, enter `date` at the `operation` level.

```
localhost::operation> date

Jan 11  16:32:01 -05:00 1996
```

This command also enables the user to set the current date and time on the switch. To set or change this information, enter the following parameters:

```
localhost::operation> date ?

localhost::operation> date [<mm/dd/yyyy hh:mm:ss> [<(+|-)hh:mm>]]
```

mm/dd/yyyy Indicates the current date. Enter the month, the day, and the year numerically. For example, enter 01/11/1996.

hh:mm:ss Indicates the current time. Enter the hour (in terms of a 24-hour clock; i.e., 1:00 pm is 13), the minutes, and the seconds. For example, to set the time as 2:02 pm, enter 14:02:00.

<(+|-)hh:mm> Indicates your time difference from Greenwich Mean Time. Enter the hours and minutes. For example, if your time zone is two and one-half hours ahead of Greenwich Mean Time, enter +02:30. Currently, changing this parameter has no effect.

NOTE: You must reboot the switch for these changes to take effect.

C.4 FLASH Operation Commands

These commands enable management of the FLASH memory system. Typing flash ? at the prompt at the operation level displays the flash commands as follows:

```
localhost::operation> flash ?

copy            delete          dir             free
get             init            put             rename
```

C.4.1 Copying a File to FLASH Memory

This command allows the user to copy a file within the FLASH memory system. Enter the following parameters:

```
localhost::operation flash> copy <from> <to>

localhost::operation flash>
```

from Indicates the file to be copied.

to Indicates the file within the FLASH memory system to which the first file will be copied.

C.4.2 Deleting a File from FLASH Memory

This command allows the user to delete a file from the FLASH memory system. Enter the following parameters:

```
localhost::operation flash> delete <file>
```

file Indicates the name of the file in FLASH memory that is to be deleted.

In order to delete a directory from the FLASH memory system (e.g., ft343.24), you must first delete all files in that directory. For example, you would list all directories in your FLASH memory system as follows:

```
localhost::operation flash> dir
FT343.24
FT343.25
CURRENT
```

Then you need to list all files in the directory that you want to delete as follows:

```
localhost::operation flash> dir ft343.24
FOREOS.EXE
```

Now delete the file in the directory as follows:

```
localhost::operation flash> del ft343.24/foreos.exe
```

Now you can delete the directory as follows:

```
localhost::operation flash> del ft343.24
```

C.4.3 Displaying the FLASH Memory Directory

This command enables the user to display the directory listing of the FLASH memory system. Enter the following parameters:

```
localhost::operation flash> dir

ft343.24
ft343.25
CURRENT
```

C.4.4 Displaying Free Space on the FLASH File

This command lets the user display the amount of remaining free space in the FLASH memory system. Enter the following parameters:

```
localhost::operation flash> free
There are 1891974 bytes of flash still available
```

C.4.5 Getting a FLASH File

This command lets the user retrieve a file from a remote host.

NOTE: The remote host must be a tftpboot server.

Enter the following parameters:

```
localhost::operation flash> get <host:remotefile> <localfile>
```

host:remotefile Indicates the name of the host and file from which the file is to be retrieved.

localfile Indicates the name of the FLASH file where the retrieved file is to be stored.

C.4.6 Initializing the FLASH File

This command lets the user initialize the FLASH file.

CAUTION

Initializing the FLASH file deletes all information from the FLASH file, including the switch software.

Because this action results in the removal of data, the switch asks you to verify this action before it re-initializes the FLASH file. Enter the following parameters:

```
localhost::operation flash> init

Are you sure you want to format the flash [n]? n

localhost::operation flash>
```

C.4.7 Putting a FLASH File on a Remote Host

This command allows the user to put a FLASH file on a remote host. Enter the following parameters:

NOTE: The remote host must be a tftpboot server.

```
localhost::operation flash> put <localfile> <host:remotefile>
```

localfile Indicates the name of the FLASH file to be copied.

host:remotefile Indicates the name of the host and file to which the FLASH file is to be copied.

C.4.8 Renaming a FLASH File

This command enables the user to rename a file that is in FLASH memory. Enter the following parameters:

```
localhost::operation flash> rename <from> <to>
```

from Indicates the current name of the file to be renamed.

to Indicates the new name of the file to be renamed.

C.5 Setting or Changing the Password

This command allows the user to set or to change the switch's administrative password.

NOTE: Be sure that this is the action you want to take because you will be prompted for a new password immediately.

If you do not type a long enough password, you will be prompted to do so. For security reasons, your keystrokes are not echoed when you enter the new password. You will be asked to verify the change by entering the new password again. If you enter a password less than five characters long, you will be asked to use a longer password. The user can get to this level by entering `password` at the operation level. Enter the following parameters:

```
localhost::operation> password

Old password:
New password:

Please use a longer password.
New password:
Retype new password:

password changed.
```

C.6 Upgrading the Switch

This command allows the user to upgrade the software on an individual SCP.

NOTE: The remote host on which the upgrade file resides must be a tftpboot server.

NOTE: To perform the initial switch software upgrade successfully, the bootp server and the tftpboot server must be configured properly. For complete instructions about performing a software upgrade, see Chapter 4, Software Upgrade Instructions, in the ATM Switch User's Manual.

The user can get to this level by entering `upgrade` at the `operation` level. Enter the following parameters:

```
localhost::operation> upgrade <remotehost>:<full path to remotefile>
```

remotehost	Indicates the IP address of the remote host on which the upgrade file resides.
full path to remotefile	Indicates the full path name of the upgrade file.

C.7 Displaying and Changing the Version of Software

This command allows the user to display and/or change the version of software that is currently running on the SCP. To display the current version, enter the following parameters:

```
localhost::operation> version
Current software version is FT340.11
Software versions installed : FT330.35 FT340.11
```

If more than one version is installed, you can type the following parameters to change the current version:

```
localhost::operation> version [<new-version>]
```

new-version Indicates the name of the software version with which the user wants to replace the current version.

NOTE: For complete instructions about changing between multiple versions of software, see Chapter 4, Software Upgrade Instructions, in the ATM Switch User's Manual.

C.8 Rebooting the Switch

This command enables the user to reboot the SCP. The user can get to this level by entering reboot at the operation level. You are asked to verify that you want to take this action. Enter the following parameters:

```
localhost::operation> reboot

Are you sure you want to reboot this switch [n]? y
```

Upon reboot, the SCP immediately closes the user out of all open AMI sessions.

APPENDIX D AMI Statistics Commands

This chapter contains a detailed description of the AMI statistics commands that display operational performance and error information for the various hardware and software features of the switch and the network modules. The main statistics menu can be found at the root level. There are several commands available under statistics. Typing statistics ? at the prompt at the root level displays the statistics commands as follows:

```
localhost::statistics> ?
  aal0              aal4              aal5              atm
  board             ctlport           ds1               ds3
  e1                e3                icmp              interface
  ip                j2                module            port
  sonet             spans             tcp               tp25
  udp               uni30             vcc               vpc
```

Each of these commands is described in the following subsections.

D.1 AAL0 Statistics

You can display AAL0 statistics for an individual switch fabric by entering aal0 at the statistics level.

```
localhost::statistics> aal0
Interface               XmtCell         RcvCell         CellDsc
asx0                          0               0               0
qaa0                          0               0               0
```

The fields in this display have the following meanings:

Interface Displays the AAL0 interface.

XmtCell Lists the number of transmitted cells.

RcvCell Shows the number of received cells.

CellDsc Displays the number of discarded cells.

D.2 AAL4 Statistics

You can display AAL4 statistics for an individual switch fabric by entering `aal4` at the `statistics` level.

```
localhost::statistics> aal4

Intfce  XmtCell  RcvCell  XmtPDU  RcvPDU  CRCErr  SARErr   CSErr  CellDsc
asx0    291691   236757   132644   58303       0       0       0       0
```

The fields in this display have the following meanings:

Interface	Displays the AAL4 interface.
XmtCell	Lists the number of transmitted cells.
RcvCell	Shows the number of received cells.
XmtPDU	Indicates the number of PDU packets transmitted.
RcvPDU	Displays the number of PDU packets received.
CRCErr	Lists the number of CRC errors.
SARErr	Shows the number of segmentation and reassembly errors.
CSErr	Indicates the number of convergence sublayer errors.
CellDsc	Displays the number of discarded cells.

D.3 AAL5 Statistics

You can display AAL5 statistics for an individual switch fabric by entering `aal5` at the `statistics` level.

```
localhost::statistics> aal5

Intfce  XmtCell  RcvCell   XmtPDU   RcvPDU  CRCErrs   CSErrs  CellDsc   PDUDsc
asx0      77683     3840    45888    16437        0        0      718     1247
qaa0    1322913   103970   719509   215460        0        0        0        0
```

The fields in this display have the following meanings:

Interface	Displays the AAL5 interface.
XmtCell	Lists the number of transmitted cells.
RcvCell	Shows the number of received cells.
XmtPDU	Indicates the number of PDU packets transmitted.
RcvPDU	Displays the number of PDU packets received.
CRCErrs	Lists the number of CRC errors.
CSErrs	Indicates the number of convergence sublayer errors.
CellDsc	Displays the number of discarded cells.
PDUDsc	Shows the number of discarded PDU packets.

D.4 ATM Statistics

You can display ATM statistics for an individual switch fabric by entering atm at the statistics level.

```
localhost::statistics> atm

Interface    XmtCell    RcvCell    VPI-OOR    VPI-Noc    VCI-OOR    VCI-Noc
asx0          412064     275218          0          0          0      11428
qaa0         6958226    5743816          0          0          0          0
```

The fields in this display have the following meanings:

Interface Displays the ATM interface.

XmtCell Lists the number of transmitted cells.

RcvCell Shows the number of received cells.

VPI-OOR Indicates the number of VPIs out of range.

VPI-Noc Lists the number of VPIs with no connection which means that there is no mapping entry listed for them.

VCI-OOR Indicates the number of VCIs out of range.

VCI-Noc Lists the number of VCIs with no connection which means that there is no mapping entry listed for them.

D.5 Switch Board Statistics

You can display switch board statistics for an individual switch board by entering `board` at the `statistics` level.

```
localhost::statistics> board

Board      VPI-Lookup-Errors      VCI-Lookup-Errors
1                       1241                    562
```

The fields in this display have the following meanings:

Board	Shows the board (switch fabric number).
VPI-Lookup-Errors	Lists the number of cells that do not match any VPI lookup tables.
VCI-Lookup-Errors	Displays the number of cells that do not match any VCI lookup tables.

D.6 Control Port Statistics

You can list the control port statistics for an individual switch fabric by entering `ctlport` at the `statistics` level.

```
localhost::statistics> ctlport

Interface        Framing-Errors      CRC-Errors
asx0                          0               0
qaa0                          0               0
```

The fields in this display have the following meanings:

Interface	Displays the control port interface.
Framing-Errors	Shows the number of framing errors.
CRC-Errors	Lists the number of CRC errors.

D.7 DS-3 Statistics

You can list statistics about all of the DS-3 network modules in an individual switch fabric by entering ds3 at the statistics level. This command is available only when at least one DS-3 network module is installed in the switch fabric.

```
localhost::statistics> ds3

ds3 Port 1C1 Counter            Value           Delta
------------------------------  --------------  --------------
ds3FramingLOSs                               0               0
ds3FramingLCVs                               0      3047081531
ds3FramingSumLCVs                   3533754755               0
ds3FramingFERRs                      122265891               0
ds3FramingOOFs                           59758               0
ds3FramingFERFs                              0               0
ds3FramingAISs                               0               0
ds3FramingPbitPERRs                 1713276195               0
ds3FramingCbitPERRs                 1583241699               0
ds3FramingFEBEs                      976095339               0
ds3PlcpFERRs                            956136               0
ds3PlcpLOFs                                  0               0
ds3PlcpBIP8s                         252419904               0
ds3PlcpFEBEs                         603317432               0
ds3PlcpYellows                               0               0
ds3AtmHCSs                           176407092               0
ds3AtmRxCells                       3709807680               0
ds3AtmTxCells                           203023               0
```

NOTE: All of the PLCP counters listed above and the Yellow counter have meaningful values only when the DS-3 network module is running in the PLCP mode; they are all meaningless when running in the HCS mode. However, the HCS counter always has meaning, regardless of which mode is running.

The fields in this display have the following meanings:

ds3FramingLOSs	Specifies the number of seconds in which Loss Of Signal (LOS) errors were detected by the DS3 Receive Framer block.
ds3FramingLCVs	Indicates the number of Line Code Violations (LCV) that were detected by the DS3 Receive Framer block.
ds3FramingSumLCVs	Shows the number of DS3 information blocks (85 bits) which contain one or more Line Code Violations (LCV).
ds3FramingFERRs	Displays the number of DS3 framing error (FERR) events.
ds3FramingOOFs	Specifies the number of DS3 Out Of Frame (OOF) error events.
ds3FramingFERFs	Indicates the number of seconds in which a Far End Receive Failure (FERF) state has been detected by the DS3 Receive Framer block. The FERF signal alerts the upstream terminal that a failure has been detected along the downstream line.
ds3FramingAISs	Shows the number of seconds in which Alarm Indication Signals (AIS) were detected by the DS3 Receive Framer block. AIS means that an upstream failure has been detected by the far end.
ds3FramingPbitPERRs	Lists the number of P-bit parity error (PERR) events.
ds3FramingCbitPERRs	Shows the number of C-bit parity error (PERR) events.
ds3FramingFEBEs	Indicates the number of DS3 far end block error (FEBE) events.
ds3PlcpFERRs	Lists the number of Physical Layer Convergence Protocol (PLCP) octet error events.
ds3PlcpLOFs	Shows the number of seconds in which Loss Of Frame (LOF) errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. LOF is declared when an Out-Of-Frame state persists for more than 1ms. LOF is removed when an in-frame state persists for more than 12ms.

ds3PlcpBIP8s Lists the number of BIP-8 (Bit Interleaved Parity - 8) error events. The BIP-8 is calculated over the Path Overhead field and the associated ATM cell of the previous frame. A BIP-N is a method of error monitoring. An N-bit code is generated by the transmitting equipment in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the previous VT SPE, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc.

ds3PlcpFEBEs Specifies the number of ATM Far End Block Error (FEBE) events.

ds3PlcpYellows Shows the number of seconds in which Yellow alarm errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. The yellow alarm is asserted when 10 consecutive yellow signal bits are set to logical 1. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

ds3AtmHCSs Lists the number of the header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

ds3AtmRxCells Shows the number of ATM cells that were received, not including idle/unassigned cells.

ds3AtmTxCells Displays the number of ATM cells that were transmitted, not including idle/unassigned cells.

D.8 E-3 Statistics

You can display statistics about all of the E-3 network modules in an individual switch fabric by entering `e3` at the `statistics` level. This command is available only when at least one E-3 network module is installed in the switch fabric.

```
localhost::statistics> e3

e3 Port 1D1 Counter                 Value           Delta
------------------------------  --------------  --------------
e3FramingLOSs                          85974               0
e3FramingLCVs                      3684415794              0
e3FramingFERRs                       85173622              0
e3FramingOOFs                          85974               0
e3FramingFERFs                             0               0
e3FramingAISs                              0               0
e3FramingBIP8s                     636877586               0
e3FramingFEBEs                       2465566               0
e3PlcpFERRs                                0               0
e3PlcpLOFs                           171950                0
e3PlcpBIP8s                                0               0
e3PlcpFEBEs                                0               0
e3PlcpYellows                              0               0
e3AtmHCSs                                  0               0
e3AtmRxCells                               0               0
e3AtmTxCells                         281929                0

Press return for more, q to quit: q
```

NOTE: All of the PLCP counters listed above and the Yellow counter have meaningful values only when the E-3 network module is running in the PLCP mode; they are all meaningless when running in the HCS mode. However, the HCS counter always has meaning, regardless of which mode is running.

The fields in this display have the following meanings:

e3FramingLOSs	Specifies the number of seconds in which Loss Of Signal (LOS) errors were detected by the E3 Receive Framer block.
e3FramingLCVs	Displays the number of Line Code Violations (LCV) that were detected by the E3 Receive Framer block.
e3FramingFERRs	Lists the number of E3 framing error (FERR) events.
e3FramingOOFs	Shows the number of E3 Out Of Frame (OOF) error events.
e3FramingFERFs	Indicates the number of Far End Receive Failures for a port configured with HCS framing. Indicates the number of Remote Alarm Indications for a port configured with PLCP framing.
e3FramingAISs	Lists the number of seconds in which Alarm Indication Signals (AIS) were detected by the E3 Receive Framer block. AIS indicates that an upstream failure has been detected by the far end.
e3FramingFEBEs	Displays the number of E3 far end block error (FEBE) events.
e3FramingBIP8s	Shows the number of E3 G.832 BIP-8 errors. This counter is only valid for a port configured with HCS framing.
e3PlcpFERRs	Lists the number of Physical Layer Convergence Protocol (PLCP) octet error events.
e3PlcpLOFs	Indicates the number of seconds in which Loss Of Frame (LOF) errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. LOF is declared when an Out-Of-Frame state persists for more than 1ms. LOF is removed when an in-frame state persists for more than 12ms.

e3PlcpBIP8s Shows the number of BIP-8 (Bit Interleaved Parity-8) error events. The BIP-8 is calculated over the Path Overhead field and the associated ATM cell of the previous frame. A BIP-N is a method of error monitoring. An N-bit code is generated by the transmitting equipment in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the previous VT SPE, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc.

e3PlcpFEBEs Displays the number of ATM Far End Block Error (FEBE) events.

e3PlcpYellows Shows the number of seconds in which Yellow alarm errors were detected by the PLCP (Physical Layer Convergence Protocol) receiver. Yellow alarm is asserted when 10 consecutive yellow signal bits are set to logical 1. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

e3AtmHCSs Lists the number of header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

e3AtmRxCells Indicates the number of ATM cells that were received, not including idle/unassigned cells.

e3AtmTxCells Displays the number of ATM cells that were transmitted, not including idle/unassigned cells.

D.9 ICMP Statistics

You can list ICMP statistics for an individual switch fabric by entering icmp at the statistics level.

```
localhost::statistics> icmp

icmp Counter                          Value              Delta
----------------------------- ------------------- --------------------
icmpInMsgs                                  815                    2
icmpInErrors                                  0                    0
icmpInDestUnreachs                           13                    0
icmpInTimeExcds                               0                    0
icmpInParmProbs                               0                    0
icmpInSrcQuenchs                              0                    0
icmpInRedirects                               0                    0
icmpInEchos                                 802                    2
icmpInEchoReps                                0                    0
icmpInTimestamps                              0                    0
icmpInTimestampReps                           0                    0
icmpInAddrMasks                               0                    0
icmpInAddrMaskReps                            0                    0
icmpOutMsgs                                 802                    2
icmpOutErrors                                 0                    0
icmpOutDestUnreachs                           0                    0
icmpOutTimeExcds                              0                    0
icmpOutParmProbs                              0                    0
icmpOutSrcQuenchs                             0                    0
icmpOutRedirects                              0                    0
Press return for more, q to quit:

icmpOutEchos                                  0                    0
icmpOutEchoReps                             802                    2
icmpOutTimestamps                             0                    0
icmpOutTimestampReps                          0                    0
icmpOutAddrMasks                              0                    0
icmpOutAddrMaskReps                           0                    0
```

The fields in this display have the following meanings:

icmpInMsgs Indicates the total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.

icmpInErrors Shows the number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).

icmpInDestUnreachs Lists the number of ICMP Destination Unreachable messages received.

icmpInTimeExcds Displays the number of ICMP Time Exceeded messages received.

icmpInParmProbs Indicates the number of ICMP Parameter Problem message received.

icmpInSrcQuenchs Shows the number of ICMP Source Quench messages received.

icmpInRedirects Displays the number of ICMP Redirect messages received.

icmpInEchos Lists the number of ICMP Echo (request) messages received.

icmpInEchoReps Indicates the number of ICMP Echo Reply messages received.

icmpInTimestamps Shows the number of ICMP Timestamp (request) messages received.

icmpInTimestampReps Displays the number of ICMP Timestamp Reply messages received.

icmpInAddrMasks Lists the number of ICMP Address Mask Request messages received.

icmpInAddrMaskReps Indicates the number of ICMP Address Mask Reply messages received.

icmpOutMsgs Shows the total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.

icmpOutErrors	Displays the number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
icmpOutDestUnreachs	Lists the number of ICMP Destination Unreachable messages sent.
icmpOutTimeExcds	Indicates the number of ICMP Time Exceeded messages sent.
icmpOutParmProbs	Shows the number of ICMP Parameter Problem messages sent.
icmpOutSrcQuenchs	Displays the number of ICMP Source Quench messages sent.
icmpOutRedirects	Lists the number of ICMP Redirect messages sent. For a host, this object is always zero, since hosts do not send redirects.
icmpOutEchos	Indicates the number of ICMP Echo (request) messages sent.
icmpOutEchoReps	Shows the number of ICMP Echo Reply messages sent.
icmpOutTimestamps	Displays the number of ICMP Timestamp (request) messages sent.
icmpOutTimestampReps	Lists the number of ICMP Timestamp Reply messages sent.
icmpOutAddrMasks	Indicates the number of ICMP Address Mask Request messages sent.
icmpOutAddrMaskReps	Shows the number of ICMP Address Mask Reply messages sent.

D.10 Interface Statistics

You can list interface statistics for an individual switch fabric by entering `interface` at the `statistics` level.

```
localhost::statistics> interface

Interface lo0 Counter                Value                Delta
---------------------------- -------------------- --------------------
ifInOctets                                1211364                20944
ifInUcastPkts                                3933                   68
ifInNUcastPkts                                  0                    0
ifInDiscards                                    0                    0
ifInErrors                                      0                    0
ifInUnknownProtos                               0                    0
ifOutOctets                               1211364                20944
ifOutUcastPkts                               3933                   68
ifOutNUcastPkts                                 0                    0
ifOutDiscards                                   0                    0
ifOutErrors                                     0                    0
ifOutQLen                                       0                    0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

ifInOctets Indicates the total number of octets received on the interface, including framing characters.

ifInUcastPkts Shows the number of subnetwork-unicast packets delivered to a higher-layer protocol.

ifInNUcastPkts Lists the number of non-unicast (i.e., subnetwork-broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol.

ifInDiscards Displays the number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.

ifInErrors	Indicates the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
ifInUnknownProtos	Shows the number of packets received via the interface which were discarded because of an unknown or unsupported protocol.
ifOutOctets	Displays the total number of octets transmitted out of the interface, including framing characters.
ifOutUcastPkts	Lists the total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
ifOutNUcastPkts	Indicates the total number of packets that higher-level protocols requested be transmitted to a non-unicast (i.e., a subnetwork-broadcast or subnetwork-multicast) address, including those that were discarded or not sent.
ifOutDiscards	Shows the number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifOutErrors	Displays the number of outbound packets that could not be transmitted because of errors.
ifOutQLen	Specifies the length of the output packet queue (in packets).

D.11 IP Statistics

You can display IP statistics for an individual switch fabric by entering `ip` at the `statistics` level.

```
localhost::statistics> ip

ip Counter                          Value               Delta
------------------------------ ------------------- --------------------
ipInReceives                                74056                  11
ipInHdrErrors                                   0                   0
ipInAddrErrors                                  0                   0
ipForwDatagrams                                 0                   0
ipInUnknownProtos                               0                   0
ipInDiscards                                    0                   0
ipInDelivers                                74056                  11
ipOutRequests                                   0                   0
ipOutDiscards                                   0                   0
ipOutNoRoutes                                   0                   0
ipReasmReqds                                    0                   0
ipReasmOKs                                      0                   0
ipReasmFails                                    0                   0
ipFragOKs                                       0                   0
ipFragFails                                     0                   0
ipFragCreates                                   0                   0
```

The fields in this display have the following meanings:

ipInReceives Shows the total number of input datagrams received from interfaces, including those received in error.

ipInHdrErrors Lists the number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.

ipInAddrErrors Shows the number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., Class E). For entities which are not IP Gateways and, therefore, do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.

ipForwDatagrams Indicates the number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter includes only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.

ipInUnknownProtos Displays the number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.

ipInDiscards Lists the number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). This counter does not include any datagrams discarded while awaiting re-assembly.

ipInDelivers Shows the total number of input datagrams successfully delivered to IP user-protocols (including ICMP).

ipOutRequests Lists the total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. This counter does not include any datagrams counted in ipForwDatagrams.

ipOutDiscards Displays the number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.

ipOutNoRoutes Lists the number of IP datagrams discarded because no route could be found to transmit them to their destination. This counter includes any packets counted in ipForwDatagrams which meet this "no-route" criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.

ipReasmReqds Shows the maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.

ipReasmOKs Indicates the number of IP datagrams successfully reassembled.

ipReasmFails Displays the number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, etc). This is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC-815) can lose track of the number of fragments by combining them as they are received.

ipFragOKs Lists the number of IP datagrams that have been successfully fragmented at this entity.

ipFragFails Shows the number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be; e.g., because their Don't Fragment flag was set.

ipFragCreates Indicates the number of IP datagram fragments that have been generated as a result of fragmentation at this entity.

D.12 Network Module Statistics

You can list network module statistics about all of the network modules in an individual switch fabric by entering `module` at the `statistics` level. The following statistics are displayed on all switches, except an SFCS-1000:

```
localhost::statistics> module

Module Priority   Status      Size  QLength Overflows
1A     0          enabled      585        0        0
1A     1          enabled      585        0        0
1B     0          enabled      585        0        0
1B     1          enabled      585        0        0
1C     0          enabled      585        0        0
1C     1          enabled      585        0        0
1D     0          enabled      585        0        0
```

The following statistics are displayed on an SFCS-1000 only:

```
localhost::statistics> module

Module Priority   Status      Size  QLength Overflows
4A     0          enabled     4096        0        0
4A     1          enabled     4096        0        0
4A     2          enabled     4096        0        0
4A     3          enabled     4096        0        0
4B     0          enabled     4096        0        0
4B     1          enabled     4096        0        0
4B     2          enabled     4096        0        0
4B     3          enabled     4096        0        0
4C     0          enabled     4096        0        0
4C     1          enabled     4096        0        0
4C     2          enabled     4096        0        0
4C     3          enabled     4096        0        0
4D     0          enabled     4096        0        0
4D     1          enabled     4096        0        0
4D     2          enabled     4096        0        0
4D     3          enabled     4096        0        0
4CTL   0          enabled     4096        0        0
```

The fields in these displays have the following meanings:

Module	Shows the number of each network module that is currently installed in the switch fabric. The 1 means that it is the first switch fabric. The letters show the position of the network module in the switch fabric.
Priority	Indicates the priority level for each network module.
Status	Shows whether the buffer is enabled or disabled.
Size	Displays the buffer size.
QLength	Lists the number of cells currently in this queue.
Overflows	Indicates the number of overflows in this queue.

D.13 Port Statistics

You can display port statistics about all of the ports on an individual switch fabric by entering port at the statistics level.

```
localhost::statistics> port

Port   Received-Cells     Transmitted-Cells      Port-Errors       Port-
                                                                 Overflows
1A1              6386                 49996                0              0
1A2                 0                 16121                0              0
1A3                 0                 16121                0              0
1A4                 0                 16121                0              0
1B1                 0                     0             9040              0
1B2                 0                     0             9040              0
1B3                 0                     0             9040              0
1B4                 0                     0             9040              0
1C1                 0                 16121                0              0
1C2                 0                 16121                0              0
1CTL           241397                120581                0              0
```

The fields in this display have the following meanings:

Port Indicates the port number.

Received-Cells Shows the number of cells received on this port.

Transmitted-Cells Shows the number of cells transmitted on this port.

Port-Errors Lists the number of seconds in which errored cells were dropped by this port.

Port-Overflows Shows the number of cells dropped on this port because of full queues.

D.14 SONET Statistics

You can display statistics about all of the SONET network modules on an individual switch fabric by entering `sonet` at the `statistics` level. This command is available only when at least one SONET network module is installed in the switch fabric.

```
localhost::statistics> sonet

sonet Port 1A1 Counter              Value           Delta
------------------------------  --------------  --------------
sonetSectionBIPs                   383833630         1584162
sonetSectionLOSs                       22103              99
sonetSectionLOFs                       23991              99
sonetLineBIPs                            532               0
sonetLineFEBEs                           355               0
sonetLineAISs                          23991              99
sonetLineFERFs                             6               0
sonetPathBIPs                            244               0
sonetPathFEBEs                           211               0
sonetPathLOPs                           1888               0
sonetPathAISs                          23991              99
sonetPathYellows                       23997              99
sonetAtmCorrectableHCSs                    0               0
sonetAtmUncorrectableHCSs                  0               0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

sonetSectionBIPs Shows the number of Section BIP-8 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. The calculated BIP-8 code is compared with the BIP-8 code extracted from the B1 byte of the following frame. Differences indicate that a section level bit error has occurred.

sonetSectionLOSs Lists the number of seconds in which Loss Of Signal (LOS) has occurred. A LOS is declared when 20 +/- 3ms of all zeros patterns is detected. LOS is cleared when two valid framing words are detected and during the intervening time no LOS condition is detected.

sonetSectionLOFs Specifies the number of seconds in which Loss Of Frame (LOF) has occurred. A LOF is declared when an out-of-frame (OOF) condition persists for 3ms. The LOF is cleared when an in-frame condition persists for 3ms. While in-frame the framing bytes (A1, A2) in each frame are compared against the expected pattern. Out-of-frame is declared when four consecutive frames containing one or more framing pattern errors have been received.

sonetLineBIPs Indicates the number of Line BIP-24 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. The calculated BIP-24 code is based on the line overhead and synchronous payload envelope (SPE) of the STS-3c stream. The line BIP-24 code is a bit interleaved parity calculation using even parity. The calculated code is compared with the BIP-24 code extracted from the B2 bytes of the following frame. Differences indicate that a line layer bit error has occurred.

sonetLineFEBEs Lists the number of line Far End Block Errors (FEBE) that have been detected since the last time the port has been reset.

sonetLineAISs Shows the number of seconds in which line Alarm Indication Signal (AIS) has occurred. A line AIS is asserted when a 111 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line AIS is removed when any pattern other than 111 is detected in these bits for five consecutive frames.

sonetLineFERFs Specifies the number of seconds in which line Far End Receive Failure (FERF) has occurred. A line FERF is asserted when a 110 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line FERF is removed when any pattern other than 110 is detected in these bits for five consecutive frames.

sonetPathBIPs Indicates the number of Path BIP-8 (Bit Interleaved Parity) errors that have been detected since the last time the port has been reset. A path BIP-8 error is detected by comparing the path BIP-8 byte (B3) extracted from the current frame, to the path BIP-8 computed for the previous frame.

sonetPathFEBEs Displays the number of path Far End Block Errors (FEBE) that have been detected since the last time the port has been reset. FEBEs are detected by extracting the 4-bit FEBE field from the path status byte (G1). The valid range for the 4-bit field is between 0000 and 1000, representing zero to eight errors. Any other value is interpreted as zero errors.

sonetPathLOPs Lists the number of seconds in which path Loss Of Pointer (LOP) has occurred. A path LOP is detected when a "normal pointer value" is not found in eight consecutive frames. The LOP is cleared when a "normal pointer value" is detected for three consecutive frames.

sonetPathAISs Indicates the number of seconds in which path Alarm Indication Signal (AIS) has occurred. A path AIS is asserted when an all-ones pattern is detected in the pointer bytes (H1 and H2) for three consecutive frames. It is cleared when a valid pointer is detected for three consecutive frames. AIS indicates that an upstream failure has been detected.

sonetPathYellows Shows the number of seconds in which path yellow alarm has occurred. A path yellow alarm is detected by extracting bit 5 of the path status byte. If bit 5 is high for ten consecutive frames, a yellow alarm is declared. A yellow alarm is cleared when bit 5 is low for ten consecutive frames. Yellow signals are used to alert upstream terminals of a downstream failure in order to initiate trunk conditioning on the failure circuit.

sonetAtmCorrectableHCSs Lists the number of correctable Header Check Sequence (HCS) error events that occurred since the port was reset. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

sonetAtmUncorrectable-HCSs Displays the number of uncorrectable Header Check Sequence (HCS) error events that occurred since the port was reset. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

D.15 SPANS Statistics

You can list SPANS statistics for an individual switch fabric by entering `spans` at the `statistics` level.

```
localhost::statistics> spans

Port 1D1 Counter                      Value              Delta
--------------------------------- -------------- --------------------
sigPathVCCs                                   0                    0
sigPathRestarts                               0                    0
sigPathCallsCompletions                       0                    0
sigPathCallsFailures                          0                    0
sigPathCallsRejections                        0                    0
sigPathSpansTransmittedMessages          193416                   76
sigPathSpansReceivedMessages                  0                    0
sigPathClsTransmittedMessages                 0                    0
sigPathClsReceivedMessages                    0                    0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

sigPathVCCs — Shows the number of VCCs on this signalling path.

sigPathRestarts — Shows the number of times this switch has lost and regained contact with the other side of the connection.

sigPathCallsCompletions — Displays the number of signalling requests that were completed.

sigPathCallsFailures — Lists the number of failed signalling calls.

sigPathCallsRejections — Shows the number of rejected requests.

sigPathSpansTransmitted-Messages — Shows the number of SPANS messages that were sent.

sigPathSpansReceived-Messages — Displays the number of SPANS messages that were received.

sigPathClsTransmitted-Messages — Lists the number of connectionless messages that were sent.

sigPathClsReceivedMes-sages — Shows the number of connectionless messages that were received.

D.16 TCP Statistics

You can display TCP statistics for an individual switch fabric by entering `tcp` at the `statistics` level.

```
localhost::statistics> tcp

tcp Counter                          Value               Delta
------------------------------ ------------------- --------------------
tcpActiveOpens                                   0                    0
tcpPassiveOpens                                 20                    0
tcpAttemptFails                                  0                    0
tcpEstabResets                                   1                    0
tcpCurrEstab                                     2                    0
tcpInSegs                                     4307                   10
tcpOutSegs                                    3290                    7
tcpRetransSegs                                   0                    0
```

The fields in this display have the following meanings:

tcpActiveOpens Shows the number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.

tcpPassiveOpens Lists the number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.

tcpAttemptFails Displays the number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.

tcpEstabResets Indicates the number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.

tcpCurrEstab Shows the number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.

tcpInSegs Lists the total number of segments received, including those received in error. This count includes segments received on currently established connections.

tcpOutSegs Displays the total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.

tcpRetransSegs Indicates the total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted octets.

D.17 TP25 Statistics

You can display statistics about all of the TP25 network modules in an individual switch fabric by entering `tp25` at the `statistics` level. The following TP25 command is available only when at least one TP25 network module is installed in the switch fabric.

```
localhost::statistics> tp25

tp25 Port 1A1 Counter               Value           Delta
------------------------------  --------------  --------------
tp25ErrorSymbol                     40452300               0
tp25AtmHCSs                                8               0
tp25AtmRxCells                         13722               0
tp25AtmTxCells                             0               0

Press return for more, q to quit:
```

The fields in this display are defined as follows:

tp25ErrorSymbol Shows the number of undefined symbols received.

tp25AtmHCSs Lists the number of header check sequence (HCS) error events. The HCS is a CRC-8 calculation over the first 4 octets of the ATM cell header.

tp25AtmRxCells Displays the number of ATM cells that were received.

tp25AtmTxCells Shows the number of ATM cells that were transmitted.

D.18 UDP Statistics

You can display UDP statistics for an individual switch fabric by entering `udp` at the `statistics` level.

```
localhost::statistics> udp

udp Counter                          Value               Delta
---------------------------- -------------------- --------------------
udpInDatagrams                                 0                    0
udpNoPorts                                     0                    0
udpInErrors                                    0                    0
udpOutDatagrams                                0
```

The fields in this display have the following meanings:

udpInDatagrams Shows the total number of UDP datagrams delivered to UDP users.

udpNoPorts Lists the total number of received UDP datagrams for which there was no application at the destination port.

udpInErrors Indicates the number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.

udpOutDatagrams Displays the total number of UDP datagrams sent from this entity.

D.19 UNI 3.0 Statistics

You can show UNI 3.0 statistics for an individual switch fabric by entering uni30 at the statistics level.

```
localhost::statistics> uni30

Port 1D1 Counter                        Value               Delta
--------------------------------- --------------- --------------------
q2931VCCs                                      0                   0
q2931Restarts                                  0                   0
q2931CallsCompletions                          0                   0
q2931CallsFailures                             0                   0
q2931CallsRejections                           0                   0
q2931TransmittedMessages                       0                   0
q2931ReceivedMessages                          0                   0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

q2931VCCs Indicates the number of Virtual Channel Connections (VCCs) on this signalling path.

q2931Restarts Displays the number of times the switch has lost and regained contact with the remote signalling entity on this path.

q2931CallsCompletions Shows the number of successfully completed calls on this signalling path.

q2931CallsFailures Lists the number of call failures on this signalling path.

q2931CallsRejections Indicates the number of connections on this signalling path that were rejected by the far end.

q2931TransmittedMessages Displays the total number of Q.2931 messages that have been transmitted over this signalling path.

q2931ReceivedMessages Shows the total number of Q.2931 messages that have been received on this signalling path.

D.20 VCC Statistics

You can display virtual channel statistics for an individual switch fabric by entering `vcc` at the `statistics` level.

```
localhost::statistics> vcc
Input          Output
Port  VPI  VCI Port  VPI  VCI  Uptime              Cells   RejectedCells
4A1     0    5 4CTL    0   34  0d:01:55            36168               0
4A1     0   14 4CTL    0   33  0d:01:55           325392               0
4A1     0   15 4CTL    0   32  0d:01:55           299002               0
4A1     0   16 4CTL    0   86  0d:01:55             3702               0
4A2     0    5 4CTL    0   37  0d:01:55                0               0
4A2     0   14 4CTL    0   36  0d:01:55                0               0
4A2     0   15 4CTL    0   35  0d:01:55                0               0
4A2     0   16 4CTL    0   87  0d:01:55                0               0
4A3     0    5 4CTL    0   40  0d:01:55                0               0
4A3     0   14 4CTL    0   39  0d:01:55                0               0
4A3     0   15 4CTL    0   38  0d:01:55                0               0
4A3     0   16 4CTL    0   88  0d:01:55                0               0
4A4     0    5 4CTL    0   43  0d:01:55                0               0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

Input Port Designates the incoming port number.

Input VPI Indicates the incoming virtual path number.

Input VCI Indicates the incoming virtual channel number.

Output Port Designates the outgoing port number.

Output VPI Shows the outgoing virtual path number.

Output VCI Indicates the outgoing virtual channel number.

Uptime Displays the length of time that this virtual channel has been in its current state.

Cells Lists the number of cells that were transmitted on this channel.

RejectedCells Lists the number of cells over this channel that were rejected (tagged or dropped) by the hardware due to a traffic policing violation.

D.21 VPC Statistics

You can display virtual path statistics for an individual switch fabric by entering `vpc` at the `statistics` level.

```
localhost::statistics> vpc
Input      Output
Port  VPI  Port  VPI  Uptime              Cells   RejectedCells
4A1     0  terminate  0d:01:56           783856               0
4A2     0  terminate  0d:01:56                0               0
4A3     0  terminate  0d:01:56                0               0
4A4     0  terminate  0d:01:56                0               0
4C1     0  terminate  0d:01:56                0               0
4C2     0  terminate  0d:01:56                0               0
4C3     0  terminate  0d:01:56            20049               0
4C4     0  terminate  0d:01:56                0               0
4D1     0  terminate  0d:00:08            14034               0
4D2     0  terminate  0d:00:08                0               0
4D3     0  terminate  0d:00:08             7232               0
4D4     0  terminate  0d:00:08                0               0
4E1     0  terminate  0d:01:56                0               0
4E2     0  terminate  0d:01:56                0               0
4E3     0  terminate  0d:01:56                0               0
4CTL    0  terminate  0d:01:56          2518080               0
originate  4A1     0  0d:01:56           778309               0

Press return for more, q to quit: q
```

The fields in this display have the following meanings:

Input Port Designates the incoming port number.

Input VPI Indicates the incoming virtual path number.

Output Port Designates the outgoing port number.

Output VPI Shows the outgoing virtual path number.

Uptime Displays the length of time that this virtual path has been in its current state.

Cells Lists the number of cells transmitted on this path.

RejectedCells Lists the number of cells over this path that were rejected (tagged or dropped) by the hardware due to a traffic policing violation.

APPENDIX E SNMP Configuration

The switch control software for the ATM switches includes an SNMP agent. The SNMP agent enables the remote monitoring and configuration of these switches.

E.1 SNMP Indexing

There are two main SNMP indexing schemes used: software port indices and hardware port indices. Software port indices are single numbers starting at 0 for the first port, incrementing 8 ports per module on a 9A000, SFCS-200WG, and on an SFCS-200BX. For example, port A1 on a 9A000, SFCS-200WG, or an SFCS-200BX has a software port index of 0. Port C3 on a 9A000, SFCS-200WG, or an SFCS-200BX has a software port index of 18, or 8 * 2 + 2 .

Hardware port indices are of the form {board}.{network module}.{port} or bnp notation. They start at 0.0.0 for the first port and increment across boards, network modules, and ports. For example, port C3 on a 9A000, SFCS-200WG, or an SFCS-200BX is hardware port 0.2.2.

Please refer to Table E.1 for a summary of the port number conventions used in Cabletron switches and related SNMP indexing format.

Table E.1 - 9A000/SFCS-200WG/SFCS-200BX

Port Name	Software Port Number	Board-Netmod-Port Index	Port Name	Software Port Number	Board-Netmod-Port Index
A1	0	0.0.0	C1	16	0.2.0
A2	1	0.0.1	C2	17	0.2.1
A3	2	0.0.2	C3	18	0.2.2
A4	3	0.0.3	C4	19	0.2.3
A5	4	0.0.4	C5	20	0.2.4
A6	5	0.0.5	C6	21	0.2.5
B1	8	0.1.0	D1	24	0.3.0
B2	9	0.1.1	D2	25	0.3.1
B3	10	0.1.2	D3	26	0.3.2
B4	11	0.1.3	D4	27	0.3.3
B5	12	0.1.4	D5	28	0.3.4
B6	13	0.1.5	D6	29	0.3.5
			CTL	56	0.7.0

E.2 SNMP Traps

SNMP traps are used to update the state of the network automatically to remote network management hosts. The SNMP agent on the switch supports several SNMP traps.

The traps generated by the switch's SNMP agent can be sent to as many destinations as needed. These destinations are configurable via the ATM Management Interface (AMI). Each destination must be an IP address of a network management host. The network management host specified for a trap destination can be any host with which the switch has connectivity. This means that the host does not have to be a directly connected ATM host. It can be on any attached network. The following table describes the supported traps.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
0	asxSwLinkDown	An asxSwLinkDown trap signifies that the sending protocol entity recognizes a failure in one of the ATM Switch links that is connected to another switch.
1	asxSwLinkUp	An asxSwLinkUp trap signifies that the sending protocol entity recognizes that one of the ATM Switch links that is connected to another switch has come up.
2	asxHostLinkDown	An asxHostLinkDown trap signifies that the sending protocol entity recognizes a failure in one of ATM Switch links that is connected to a host.
3	asxHostLinkUp	An asxHostLinkUp trap signifies that the sending protocol entity recognizes that one of the ATM Switch links that is connected to a host has come up.
4	asxNetModuleDown	An asxNetModuleDown trap signifies that the sending protocol entity recognizes a failure in one of the ATM Switch network modules, that is identified by the board and the module numbers. This is probably caused by a hot-swap of a network module.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
5	asxNetModuleUp	An asxNetModuleUp trap signifies that the sending protocol entity recognizes a new operational ATM Switch network module, that is identified by the board and the module numbers. This is probably caused by a hot-swap of a network module.
6	asxPsInputDown	This trap alerts that one ATM switch power supply failed due to failure in the input voltage. The power supply that failed is identified by the power supply index. Note that an input voltage may be out of specification and may not cause a power supply failure if high loads are not applied.
7	asxPsInputUp	This trap alerts that one ATM switch power supply that had an AC input failure is up. The power supply that is back up is identified by the power supply index.
9	asxPsOutputDown	This trap alerts that one ATM switch power supply output or the power supply was physically removed. The power supply that failed is identified by the power supply index.
10	asxPsOutputUp	This trap alerts that one ATM switch power supply that had an output failure or was removed is now up. The power supply that is back up is identified by the power supply index.
22	asxFanBankDown	This trap alerts that one ATM switch fan bank failed. The fan bank that failed is identified by the fan bank index.
23	asxFanBankUp	This trap alerts that one ATM switch fan bank is up. The fan bank that is back up is identified by the fan bank index.
28	asxLinkDown	This trap alerts that the link that is identified by {hwPortBoard, hwPortModule, hwPortNumber} was configured up but lost its carrier (or the framing bit) and is currently down.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
29	asxLinkUp	This trap alerts that the link that is identified by {hwPortBoard, hwPortModule, hwPortNumber} is back up.
30	asxSpansDown	This trap alerts that the SPANS signalling on the link that is identified by the sigPathPort and sigPathVPI failed.
31	asxSpansUp	This trap alerts that the SPANS signalling on the link that is identified by the sigPathPort and sigPathVPI is up.
32	asxTempSensorOverTemp	This trap alerts that one of the temperature sensors indicates over temperature. The temperature sensor is identified by the temperature sensor index.
33	asxTempSensorRegularTemp	This trap alerts that one of the temperature sensors indicates regular temperature. The temperature sensor is identified by the temperature sensor index.
34	asxFabricTemperature-OverTemp	This trap alerts that one of the temperature sensors indicates over temperature. The temperature sensor is identified by the temperature sensor index.
35	asxFabricTemperature-RegularTemp	This trap alerts that one of the temperature sensors indicates regular temperature. The temperature sensor is identified by the temperature sensor index.
36	asxSonetLOSon	This trap indicates that the specified SONET port is experiencing Loss Of Signal. Bellcore Document TA-NWT-000253 Section 6.3.1.1.1 states that... A SONET NE shall declare a LOS failure when the LOS defect persists for 2.5 (+- .5) seconds, or when a LOS defect is present and the criteria for LOF failure declaration have been met.
37	asxSonetLOSoff	This trap indicates that the LOS condition identified by trap asxSonetLOSon has been cleared.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
38	asxSonetPathLabelOn	This trap indicates that the specified SONET port is receiving and errored C2 Path Label byte. Reference Bellcore Document TA-NWT-000253 Section 3.3.2.3 and 6.3.1.1.8 the Path Label (C2) byte should have the value 0x13.
39	asxSonetPathLabelOff	This trap indicates that the Errored Path Label (C2) byte error condition signalled by the asxSonetPathLabelOn trap has been cleared.
40	asxSonetLineAISon	This trap indicates that the specified SONET port is receiving a Line level Alarm Indication Signal from the far-end equipment.
41	asxSonetLineAISoff	This trap indicates that the Line AIS error condition signalled by the asxSonetLineAISon trap has been cleared.
42	asxDS3FERFOn	This trap indicates that the specified DS3 port is in the DS3 Yellow Alarm or FERF state. The FERF or DS3 Yellow alarm is declared if either OOF(LOF), LOS or AIS is detected and persists for 2.5+- .5 seconds.
43	asxDS3FERFOff	This trap indicates that the specified DS3 port is no longer in the FERF or DS3 Yellow Alarm state.
44	asxDS3PLCPYellowOn	This trap indicates that the specified DS3 port is in the PLCP Yellow Alarm state. The Yellow alarm is declared if PLCP LOF is detected and persists for 2.5+- .5 seconds.
45	asxDS3PLCPYellowOff	This trap indicates that the specified DS3 port is no longer in the PLCP Yellow Alarm state.
46	asxDS3PLCPYellowDetected	This trap indicates that the specified DS3 port has detected incoming Yellow Alarm.
47	asxDS3PLCPYellowCleared	This trap indicates that the specified DS3 port has detected clearance of incoming Yellow Alarm.
48	asxDS3PLCPLOFDetected	This trap indicates that the specified DS3 port has detected incoming LOF Alarm.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
49	asxDS3PLCPLOFCleared	This trap indicates that the specified DS3 port has detected clearance of incoming LOF Alarm.
50	asxDS3LOFDetected	This trap indicates that Loss Of Frame(LOF) is detected on the incoming signal.
51	asxDS3LOFCleared	This trap indicates that Loss Of Frame is cleared on the incoming signal.
52	asxDS3AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.
53	asxDS3AISCleared	This trap indicates that AIS Alarm is cleared on the incoming signal.
54	asxDS3LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.
55	asxDS3LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.
56	asxDS1YellowOn	This trap indicates that the specified DS1 port is in the Yellow Alarm state. The Yellow alarm is declared if either OOF or AIS is detected and persists for 2.5+- .5 seconds.
57	asxDS1YellowOff	This trap indicates that the specified DS1 port is no longer in the Yellow Alarm state.
58	asxDS1PLCPYellowOn	This trap indicates that the specified DS1 port is in the PLCP Yellow Alarm state. The Yellow alarm is declared if PLCP LOF is detected and persists for 2.5+- .5 seconds.
59	asxDS1PLCPYellowOff	This trap indicates that the specified DS1 port is no longer in the PLCP Yellow Alarm state.
60	asxDS1PLCPYellowDetected	This trap indicates that the specified DS1 port has detected an incoming Yellow Alarm.
61	asxDS1PLCPYellowCleared	This trap indicates that the specified DS1 port has detected clearance of an incoming Yellow Alarm.
62	asxDS1PLCPLOFDetected	This trap indicates that the specified DS1 port has detected an incoming LOF Alarm.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
63	asxDS1PLCPLOFCleared	This trap indicates that the specified DS1 port has detected clearance of an incoming LOF Alarm.
64	asxDS1YellowDetected	This trap indicates that Yellow Alarm is detected on the incoming signal.
65	asxDS1YellowCleared	This trap indicates that Yellow Alarm is cleared on the incoming signal.
66	asxDS1AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.
67	asxDS1AISCleared	This trap indicates that AIS Alarm is cleared on the incoming signal.
68	asxDS1LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.
69	asxDS1LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.
70	asxDS1LOFDetected	This trap indicates that LOF Alarm is detected on the incoming signal.
71	asxDS1LOFCleared	This trap indicates that LOF Alarm is cleared on the incoming signal.
74	asxDS3FERFDetected	This trap indicates that FERF Alarm is detected on the incoming signal.
75	asxDS3FERFCleared	This trap indicates that FERF Alarm is cleared on the incoming signal.
110	asxJ2YellowOn	This trap indicates that the specified J2 port is in the Yellow Alarm state. The Yellow alarm is declared if either LOF or LOS or AIS is detected and persists for 2.5+- .5 seconds.
111	asxJ2YellowOff	This trap indicates that the specified J2 port is no longer in the Yellow Alarm state.
112	asxJ2YellowDetected	This trap indicates that Yellow Alarm is detected on the incoming signal.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
113	asxJ2YellowCleared	This trap indicates that Yellow Alarm is cleared on the incoming signal.
114	asxJ2AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.
115	asxJ2AISCleared	This trap indicates that AIS Alarm is cleared on the incoming signal.
116	asxJ2LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.
117	asxJ2LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.
118	asxJ2LOFDetected	This trap indicates that LOF Alarm is detected on the incoming signal.
119	asxJ2LOFCleared	This trap indicates that LOF Alarm is cleared on the incoming signal.
1024	asxOutputQueueCongested	This trap indicates that the output queue for the given priority has exceeded its dedicated length, and has begun overflowing into the shared buffer space on the network module.
1025	asxOutputQueueCellLoss	This trap indicates that the output queue for the given priority has overflowed and cells have been dropped.
1026	asxExtendedModeViolation	This trap indicates that a series A or B network module was inserted into a switch board running in extended mode.
1027	asxNonextendedMode-Warning	This trap indicates that a series C or greater network module was inserted into a switch board running in non-extended mode.
1028	q2931AVRejectTrap	This trap is generated whenever any UNI3.x with AddressValidation enabled rejects a Setup Request call more than q2931AVRejectTrapThreshold times in any given q2931AVRejectTrapPeriod.

Table E.2 - SNMP Traps Supported on the SFCS Switches

Trap Number	Trap Name	Description
1029	crConfMemoryOflow	This trap is generated when the allocated call record memory (as indicated by crMemoryAllocated) is exceeded.
1030	crXfrPrimaryXfrFailed	This trap is generated when the call record transfer to the primary host (as indicated by crXfrPrimaryUrl) fails.
1031	crXfrSecondaryXfrFailed	This trap is generated when the call record transfer to the secondary host (as indicated by crXfrSecondaryUrl) fails.
1032	crConfMemAllocFail	This trap is generated when Callrecord functionality is unable to allocate memory as specified by crMemoryAllocated. This can happen when the crConfAdminStatus changes state from "off" or when the switch reboots when Callrecords is configured "on".
1033	crGeneralFailure	This trap is generated when any of the callrecord related functionality fails for any reason. One example would be when the Callrecord Module fails to schedule an interval timer.

E.2.1 Adding SNMP Trap Destinations

To create one or more SNMP trap destinations on a Cabletron switch, log in to AMI and open a session on the switch. Enter the following parameters:

configuration snmp trap new <ipaddress>

The <ipaddress> variable indicates the IP address of the SNMP trap destination that is to be created. Repeat this for as many SNMP trap destinations as needed. Traps are active as soon as they are set.

E.2.2 Displaying SNMP Trap Destinations

To list all of the current SNMP trap destinations, log in to AMI and open a session on the switch. The SNMP traps supported by this switch are detailed in the FORE-Switch-MIB. Enter the following parameters:

configuration snmp trap show

```
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
5       198.29.31.130
```

E.2.3 Removing SNMP Trap Destinations

To delete one or more SNMP trap destinations for a Cabletron switch, log in to AMI and open a session on the switch. Prior to deleting any trap that may need to be recreated later, as a precaution, a recommended practice is to list all trap destinations using AMI and either copy the screen or write down the destinations. To delete a trap, enter the following parameters:

configuration snmp trap delete <trap>

The <trap> variable indicates the IP address of the SNMP trap destination that is to be removed. Repeat this for as many SNMP trap destinations as needed.

APPENDIX F *ForeThought* PNNI

PNNI (Private Network Node Interface or Private Network-to-Network Interface) is a protocol defining interoperability between private ATM switches. PNNI defines both the routing and signalling standards for inter-switch interoperability.

This appendix provides a broad overview of Cabletron Systems' version of PNNI, *ForeThought* PNNI (FT-PNNI), and its use in a multiple-switch network. FT-PNNI is a scalable routing and signalling protocol used in networks containing multiple Cabletron switches. FT-PNNI simplifies large network topologies by organizing the nodes (switches) in that network into smaller groups.

It is this reorganization of the network topology that makes FT-PNNI's simplified routing possible. By segmenting a large network into smaller groups of nodes, FT-PNNI reduces the amount of topology information that those very nodes must maintain.

To understand FT-PNNI, it is important to first understand how a large network topology can be simplified by grouping the nodes contained therein. Once the organization of the physical network is determined, FT-PNNI's routing and signalling protocols can be better understood.

This appendix examines the routing and signalling that FT-PNNI utilizes. It also looks at the physical network and how it can be simplified through the use of peer groups.

F.1 FT-PNNI Routing

The FT-PNNI routing protocol serves to distribute topology and address reachability information between switches and groups of switches in a network. This topology and addressing information is used by switches to compute paths through the network. FT-PNNI routing has the following aspects:

- Hello Protocol
- Topology Database Exchange
- Flooding
- Hierarchical Routing

F.1.1 Hello Protocol

At regular intervals, each switch transmits a hello indication on each of its FT-PNNI routing channels. The time between these hello indications is called the Hello Indication Interval. When a switch receives a hello indication from one of its neighbors, it initializes the logical link (loglink) from that neighbor to itself in the topology database.

Since hello indications are sent periodically, loglinks are also refreshed periodically. Note that loglinks discovered as a result of hello indications are unidirectional. Each switch initializes unidirectional loglinks that have an adjacent switch as the destination.

F.1.2 Topology Database Exchange

At regular intervals, called the NSAP map indication interval, each switch sends to each of its neighbors a group of loglinks taken from its topology database. This is done to ensure that the topology databases of the switches stay synchronized.

F.1.3 Flooding

Whenever a new loglink is discovered by a switch, the switch immediately broadcasts the existence of this link to all of its neighbors. The neighboring switches then broadcast the existence of the same link to all of their neighbors. Very quickly, the existence of the new loglink is flooded immediately to all of the switches in the network.

Similarly, when a link goes down, that information is propagated through the network immediately. Also, when a *significant* change is seen in the metrics of a loglink, the latest state of the loglink is propagated immediately throughout the network.

F.1.4 Hierarchical Routing

Hierarchical routing is a method of routing that uses hierarchical addressing of nodes (switches), thereby enabling the logical partitioning of networks, topology information hiding between the partitions, and increased scalability of the routing protocol. FT-PNNI supports hierarchical node identifiers and, therefore, the hierarchical organization of networks.

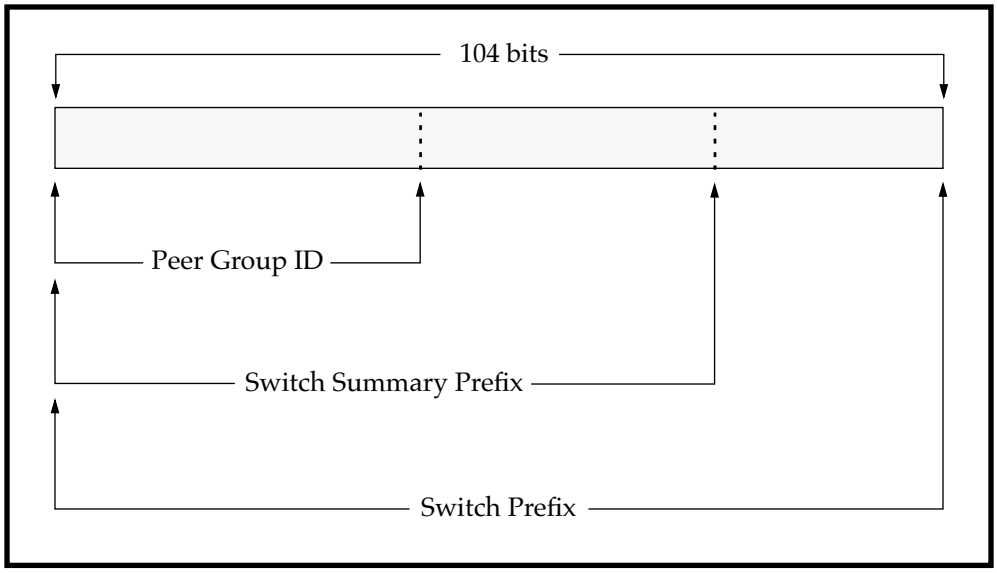

Figure F.1 - Example of a 13-byte Switch Prefix

F.1.4.1 Hierarchical Addressing

FT-PNNI uses private ATM address prefixes (NSAP prefixes) as node identifiers. FT-PNNI does not distinguish between node identifiers and reachability information. Thus, the IDs of nodes in the FT-PNNI addressing map are NSAP prefixes. In the default case at the lowest level peer group, the switches have a 13-byte prefix as their node ID and end systems (hosts) have a 19-byte prefix as their node ID.

F.1.4.1.1 Switch Prefix

Each switch in a FT-PNNI network is configured with a 13-byte prefix called the switch prefix. Hosts that are attached to the switch are presented with this prefix during ILMI address registration. It is in this way that end systems are configured with a private ATM address that includes the 13-byte switch prefix.

F.1.4.1.2 Switch Summary Prefix

Each switch is configured with a switch mask (*swmask*) which gives the length of the switch summary prefix within the switch prefix. The *swmask* gives the number of most significant bits of the switch prefix that constitute the switch summary prefix. Since all end system addresses attached to a switch have the same switch summary prefix, their reachability information can be summarized by this prefix (i.e., by the switch summary prefix).

F.1.4.1.3 Peer Group ID

Each switch is configured with a peer group mask (*pgmask*) which gives the length of the peer group ID within the switch summary prefix. The *pgmask* gives the number of most significant bits of the switch summary prefix that constitute the peer group ID.

The peer group ID is a unique identifier for every node (switch or end system) in a particular peer group. Every switch and host in peer group A, for example, will have the same peer group ID.

A simple example of summarizing by peer group ID can be seen in Figure F.3, where every switch (and end system, although not shown) in peer group A is identified starting with "A."

F.1.4.2 Path Computation

Path computation is performed on demand whenever FT-PNNI signalling requests a path to a given destination. The Bellman-Ford Shortest Path algorithm is used to compute the shortest path tree of all nodes in the topology with the local node as the source. The administrative weight metric in the loglinks is used as the minimizing criterion in computing the shortest path route. In the case of finding multiple equal-cost paths to a given destination, available cell rate is used to break ties.

F.2 The Physical Network

In an ATM network, data is sent and received over virtual circuits, or circuits that only exist when needed. This communication over these virtual circuits is made possible by signalling that occurs between the switches in the network.

In a network of Cabletron switches, any new addition to the topology is recognized immediately by all nodes (switches) having a direct connection to the new node. Then each switch that has recognized the new switch sends a message to each of its direct connections, and so on. Eventually, and within a very brief period of time, every switch in the network is aware of the new addition and the links by which that new addition can be reached. This topology is stored by each switch in a topology database.

In a small, local area network (LAN), the topology is relatively simple, meaning that the switches in the LAN have a relatively small topology database to maintain. As the LAN grows, however, and more switches are added, that same database can grow to be very large in a short period of time.

As this topology database grows, the amount of information a switch must consult when searching for an address also grows. In the end, this can result in delayed connection set-up, congestion in the network, and even lost data.

Figure F.2 depicts a typical ATM network, containing 21 Cabletron switches (⊠). The hierarchy of this network is flat, meaning that each switch must be aware of all the other switches in the network, as well as all the possible routes to those switches. As more switches are added to this network, the hierarchy will become more complex and the switches will have to contend with larger topology databases.

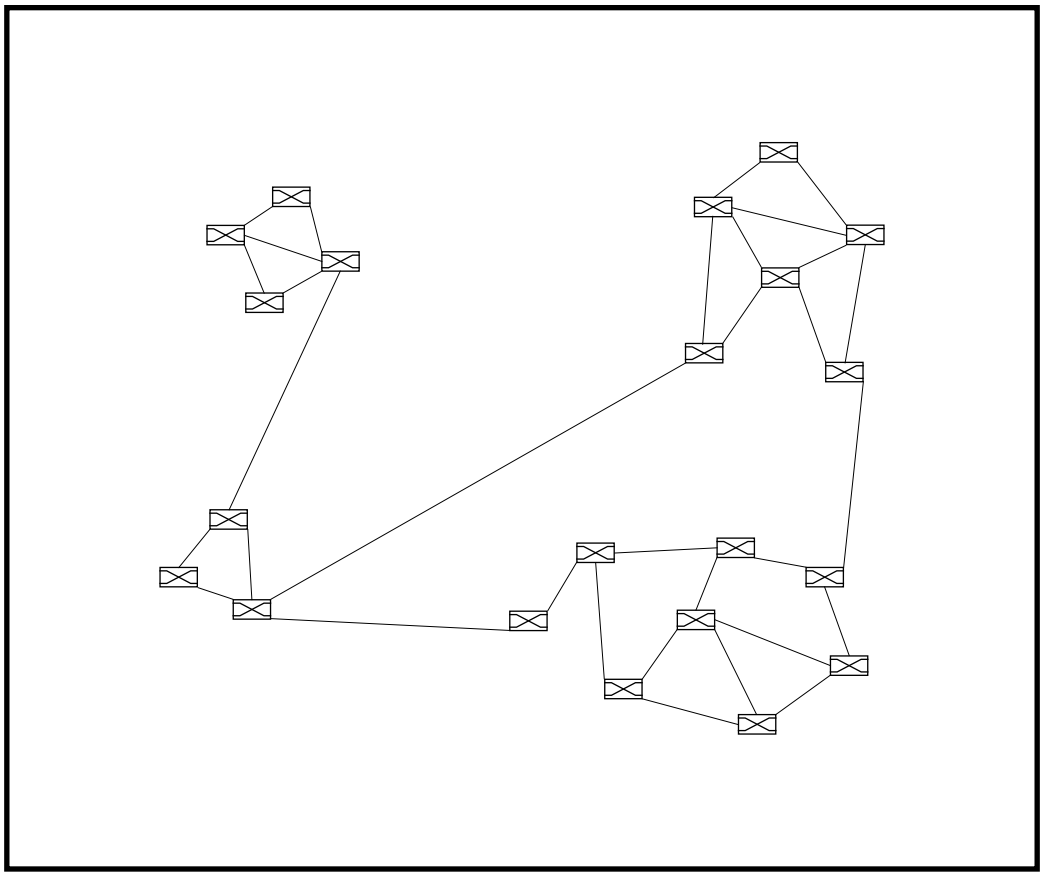

Figure F.2 - Private ATM Network with 21 Switches and 34 Bidirectional Links

It is in these large, single-level networks that FT-PNNI is most useful, because it lets you simplify large network topologies by creating a two-level hierarchy. In this hierarchy, communities of contiguous switches are grouped together and they are collectively summarized by a single, logical node.

Figure F.3 shows the same network as in Figure F.2 after being organized into peer groups, now having a two-level hierarchy. The subsections that follow explain the organization of these peer groups, how they simplify the overall network topology, and how they change the logical view of the network.

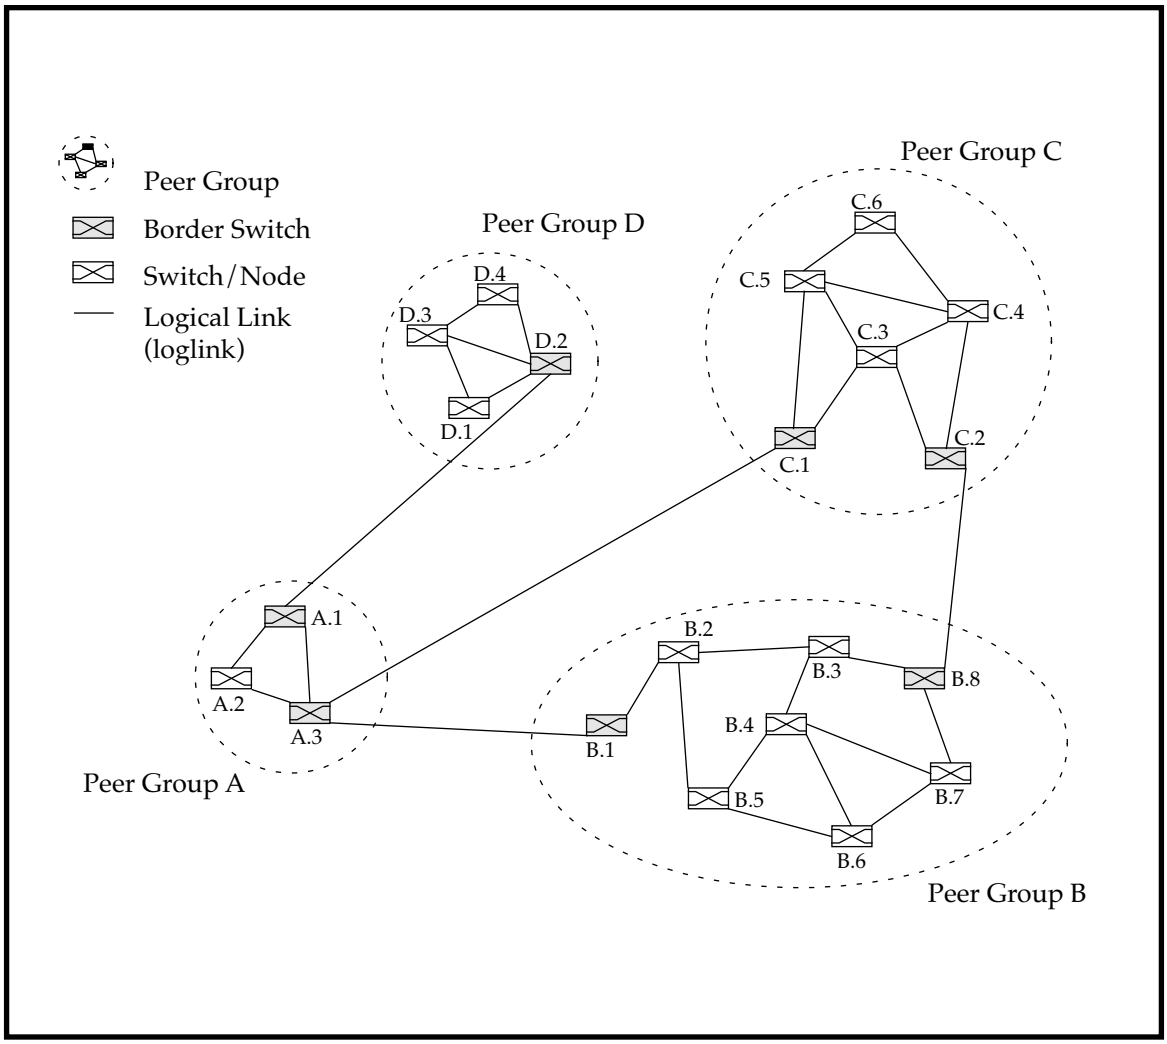

Figure F.3 - Example of FT-PNNI Hierarchy Showing Lowest-Level Peer Groups

F.2.1 Peer Groups

The FT-PNNI hierarchy begins with a network of switches, organized into peer groups. A peer group is a collection of interconnected switches that are organized into a group. Peer group organization can be determined by a network administrator, but switches that are located close to one another are usually made into a peer group.

The network shown in Figure F.3 is organized into four lowest-level peer groups: A, B, C, and D. The switches within a certain peer group are numbered according to that particular group. For example, the switches in peer group A are identified as A.1, A.2, and A.3.

Peer groups have a peer group identifier (ID), assigned at configuration time and exchanged in hello indication messages. Switches can determine in which peer group they are located by comparing these peer group IDs.

Switches in a peer group are aware of the topology of their own peer group and the existence of all other peer groups. They recognize the links between their peer group and others, but they are not aware of internal topology information of other peer groups. Instead, the topologies of other peer groups are summarized as a single, reachable location, known as a peer group summary node (PGSN).

F.2.2 Peer Group Topology

To maintain an accurate and updated view of its relative location and status, a switch sends hello indication messages to every other switch with which it has a direct connection. These hello indications contain the switch prefix, peer group membership information, and link metrics (attributes) for the physical link between the two switches.

Through this exchange of messages, each switch learns which switches are its immediate neighbors, to what peer groups they belong, and whether or not the link between itself and its neighbors is valid.

F.2.3 Border Switches

A border switch is any switch that has at least one link to a switch in another peer group. Border switches play an important role in FT-PNNI because they are responsible for summarizing reachability information for their respective peer groups, appropriately filtering the flow of topology database information across peer group boundaries, and building the lowest level source route for call setups entering the peer group.

F.2.4 Peer Group Summary Node (PGSN)

A PGSN is a hypothetical or imaginary node that summarizes a peer group's reachability information. The PGSN has the peer group ID of its peer group as its switch summary prefix. Each border switch in the peer group advertises a logical link (loglink) to the PGSN. The PGSN is a logical representation of the switches contained in a peer group.

F.2.5 Backbone Topology

Loglinks between border switches and loglinks from border switches to PGSNs are called backbone links and considered part of the backbone topology. Information regarding these backbone links is propagated across peer group boundaries during database exchange and flooding.

F.2.6 Single Switch Perspective

The main reason for grouping switches in large networks is to simplify each individual switch's view of the topology. For example, each switch in peer group A is aware of every other switch in peer group A, the border switches in the rest of the network, the links between them, and the backbone topology. Switches in peer group A are not aware, however, of the internal topology of other peer groups. Instead, individual switches see a PGSN (see Figure F.4).

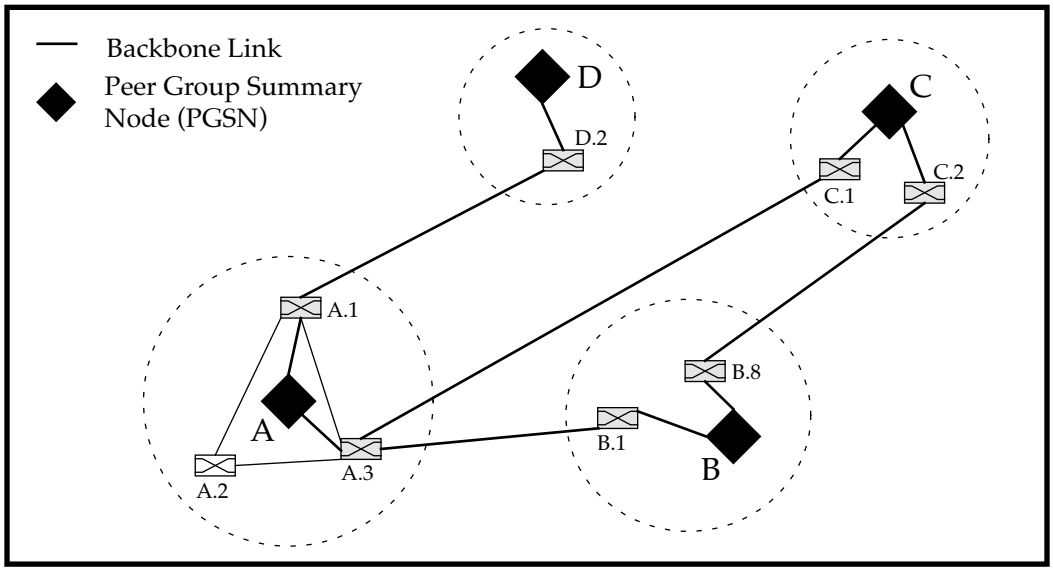

Figure F.4 - View of the Network from Switches in Peer Group A

GLOSSARY

802.1d Spanning Tree Bridging - the IEEE standard for bridging; a MAC layer standard for transparently connecting two or more LANs (often called subnetworks) that are running the same protocols and cabling. This arrangement creates an extended network, in which any two workstations on the linked LANs can share data.

802.3 Ethernet - the IEEE standard for Ethernet; a physical-layer standard that uses the CSMA/CD access method on a bus-topology LAN.

802.5 Token Ring - the IEEE physical-layer standard that uses the token-passing access method on a ring-topology LAN.

AAL (ATM Adaptation Layer) - the AAL divides the user information into segments suitable for packaging into a series of ATM cells. There are several types of AALs in use. Cabletron Systems currently supports AAL 5 and AAL 3/4. AAL 3/4 supports connection-oriented VBR data transfer and connectionless VBR data transfer, respectively. AAL 5 is defined as Simple and Efficient Adaptation Layer (SEAL).

AAL Connection - an association established by the AAL between two or more next higher layer entities.

ABR (Available Bit Rate) - a type of traffic for which the ATM network attempts to meet that traffic's bandwidth requirements. It does not guarantee a specific amount of bandwidth and the end station must retransmit any information that did not reach the far end.

Address Mask - a bit mask used to identify which bits in an address (usually an IP address) are network significant, subnet significant, and host significant portions of the complete address. This mask is also known as the subnet mask because the subnetwork portion of the address can be determined by comparing the binary version of the mask to an IP address in that subnet. The mask holds the same number of bits as the protocol address it references.

Agent (SNMP) - a component of network- and desktop-management software, such as SNMP, that gathers information from MIBs.

AIS (Alarm Indication Signal) - a line AIS is asserted when a 111 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line AIS is removed when any pattern other than 111 is detected in these bits for five consecutive frames.

AMI (ATM Management Interface) - the user interface to Cabletron Systems' switch control software (SCS). AMI lets users monitor and change various operating configurations of Cabletron Systems switches and network module hardware and software, IP connectivity, and SNMP network management.

ANSI (American National Standards Institute) - a private organization that coordinates the setting and approval of some U.S. standards. It also represents the United States to the International Standards Organization.

API (Application Program Interface) - a language format that defines how a program can be made to interact with another program, service, or other software; it allows users to develop custom interfaces with Cabletron products.

APP (application program) - a complete, self-contained program that performs a specific function directly for the user.

AppleTalk - a networking protocol developed by Apple Computer for communication between Apple's products and other computers. Independent of the network layer, AppleTalk runs on LocalTalk, EtherTalk and TokenTalk.

ARP (Address Resolution Protocol) - a method used to resolve higher level protocol addressing (such as IP) into the appropriate header data required for ATM; i.e., port, VPI, and VCI; also defines the AAL type to be used.

ASCII (American Standard Code for Information Interchange) - a standard character set that (typically) assigns a 7-bit sequence to each letter, number, and selected control characters.

Assigned Cell - a cell that provides a service to an upper layer entity or ATM Layer Management entity (ATMM-entity).

asxmon - a FORE program that repeatedly displays the state of the switch and of all its active ports.

Asynchronous time division multiplexing - a multiplexing technique in which a transmission capability is organized into a priori, unassigned time slots. The time slots are assigned to cells upon request of each application's instantaneous real need.

ATM (Asynchronous Transfer Mode) - a transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

ATM Forum - an international non-profit organization formed with the objective of accelerating the use of ATM products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness.

ATM Layer link - a section of an ATM Layer connection between two adjacent active ATM Layer entities (ATM-entities).

ATM Link - a virtual path link (VPL) or a virtual channel link (VCL).

ATM Peer-to-Peer Connection - a virtual channel connection (VCC) or a virtual path connection (VPC) directly established, such as workstation-to-workstation. This setup is not commonly used in networks.

ATM Traffic Descriptor - a generic list of parameters that can be used to capture the intrinsic traffic characteristics of a requested ATM connection.

ATM User-to-User Connection - an association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher layer entities or between two or more ATM entities). The communication over an ATM Layer connection may be either bidirectional or unidirectional. The same Virtual Channel Identifier (VCI) is used for both directions of a connection at an interface.

atmarp - a FORE program that shows and manipulates ATM ARP entries maintained by the given device driver. This is also used to establish PVC connections.

atmconfig - a FORE program used to enable or disable SPANS signalling.

atmstat - a FORE program that shows statistics gathered about a given adapter card by the device driver. These statistics include ATM layer and ATM adaptation layer cell and error counts. This can also be used to query other hosts via SNMP.

Backbone - the main connectivity device of a distributed system. All systems that have connectivity to the backbone connect to each other. This does not stop systems from setting up private arrangements with each other to bypass the backbone for cost, performance, or security.

Bandwidth - usually identifies the capacity or amount of data that can be sent through a given circuit; may be user-specified in a PVC.

BGP (Border Gateway) Protocol - used by gateways in an internet, connecting autonomous networks. It is derived from experiences learned using the EGP.

BIP (Bit Interleaved Parity) - an error-detection technique in which character bit patterns are forced into parity, so that the total number of one bits is always odd or always even. This is accomplished by the addition of a one or zero bit to each byte, as the byte is transmitted; at the other end of the transmission, the receiving device verifies the parity (odd or even) and the accuracy of the transmission.

B-ISDN (Broadband Integrated Services Digital Network) - a common digital network suitable for voice, video, and high-speed data services running at rates beginning at 155 Mbps.

Bridge - a device that expands a Local Area Network by forwarding frames between data link layers associated with two separate cables, usually carrying a common protocol. Bridges can usually be made to filter certain packets (to forward only certain traffic).

Broadband - a service or system requiring transmission channels capable of supporting rates greater than the Integrated Services Digital Network (ISDN) primary rate.

Broadband Access - an ISDN access capable of supporting one or more broadband services.

Brouter (bridging/router) - a device that routes some protocols and bridges others based on configuration information.

BUS (Broadcast and Unknown Server) - in an emulated LAN, the BUS is responsible for accepting broadcast, multicast, and unknown unicast packets from the LECs to the broadcast MAC address (FFFFFFFFFFFF) via dedicated point-to-point connections, and forwarding the packets to all of the members of the ELAN using a single point-to-multipoint connection.

CAC (Connection Admission Control) - the procedure used to decide if a request for an ATM connection can be accepted based on the attributes of both the requested connection and the existing connections.

Call - an association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.

CBR (Constant Bit Rate) - a type of traffic that requires a continuous, specific amount of bandwidth over the ATM network (e.g., digital information such as video and digitized voice).

cchan - a FORE program used to manage virtual channels on a FORE Systems ATM switch running asxd.

CCITT (International Telephone and Telegraph Consultative Committee) - the international standards body for telecommunications.

CDV (Cell Delay Variation) - a quantification of cell clumping for a connection. The cell clumping CDV (y_k) is defined as the difference between a cell's expected reference arrival time (c_k) and its actual arrival time (a_k). The expected reference arrival time (c_k) of cell k of a specific connection is $\max\ [c_{\{k-1\}} + T, a_k]$. T is the reciprocal of the negotiated peak cell rate.

CE (Connection Endpoint) - a terminator at one end of a layer connection within a SAP.

CEI (Connection Endpoint Identifier) - an identifier of a CE that can be used to identify the connection at a SAP.

Cell - an ATM Layer protocol data unit (PDU).

Cell Header - ATM Layer protocol control information.

Cell Transfer Delay - the transit delay of an ATM cell successfully passed between two designated boundaries.

CLP (Cell Loss Priority) - the last bit of byte four in an ATM cell header; indicates the eligibility of the cell for discard by the network under congested conditions. If the bit is set to 1, the cell may be discarded by the network depending on traffic conditions.

Concentrator - a communications device that offers the ability to concentrate many lower-speed channels into and out of one or more high-speed channels.

Connection - the concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

Connectionless Service - a type of service in which no pre-determined path or link has been established for transfer of information, supported by AAL 4.

Connection-Oriented Service - a type of service in which information always traverses the same pre-established path or link between two points, supported by AAL 3.

Corresponding Entities - peer entities with a lower layer connection among them.

cpath - a FORE program used to manage virtual paths on a FORE Systems ATM switch running asxd.

cport - a FORE program used to monitor and change the state of ports on a FORE Systems ATM switch running asxd.

CRC (Cyclic Redundancy Check) - an error detection scheme in which a number is derived from the data that will be transmitted. By recalculating the CRC at the remote end and comparing it to the value originally transmitted, the receiving node can detect errors.

CS (Convergence Sublayer) - a portion of the AAL. Data is passed first to the CS where it is divided into rational, fixed-length packets or PDUs (Protocol Data Units). For example, AAL 4 processes user data into blocks that are a maximum of 64 kbytes long.

DARPA (Defense Advanced Research Projects Agency) - the US government agency that funded the ARPANET.

DCS (Digital Cross-connect System) - an electronic patch panel used to route digital signals in a central office.

Demultiplexing - a function performed by a layer entity that identifies and separates SDUs from a single connection to more than one connection. (See multiplexing.)

DIP Switch (Dual In-line Package) - a device that has two parallel rows of contacts that let the user switch electrical current through a pair of those contacts to on or off. They are used to reconfigure components and peripherals.

Domain Name Server - a computer that converts names to their corresponding Internet numbers. It allows users to telnet or FTP to the name instead of the number.

DNS (Domain Name System) - the distributed name and address mechanism used in the Internet.

DSn (Digital Standard n (0, 1, 1C, 2, and 3)) - a method that defines the rate and format of digital hierarchy.

Asynchronous data rates are defined as follows:

DS0	64kb/s	1 voice channel
DS1	1.544Mb/s	24 DS0s
DS1C	3.152 Mb/s	2 DS1s
DS2	6.312 Mb/s	4 DS1s
DS3	44.736 Mb/s	28 DS1s

Synchronous data rates (SONET) are defined as:

STS-1/OC-1	51.84 Mb/s	28 DS1s or 1 DS3
STS-3/OC-3	155.52 Mb/s	3 STS-1s byte interleaved
STS-3c/OC-3c	155.52 Mb/s	Concatenated, indivisible payload
STS-12/OC-12	622.08 Mb/s	12 STS-1s, 4 STS-3cs, or any mixture
STS-12c/OC-12c	622.08 Mb/s	Concatenated, indivisible payload
STS-48/OC-48	2488.32 Mb/s	48 STS-1s, 16 STS-3cs, or any mixture

EGP (Exterior Gateway) Protocol - used by gateways in an internet, connecting autonomous networks.

EISA (Extended Industry Standard Architecture) - a bus architecture for desktop computers that provides a 32-bit data passage while maintaining compatibility with the ISA or AT architecture.

elarp - a FORE program that shows and manipulates MAC and ATM address mappings for LAN Emulation Clients (LECs).

elconfig - a FORE program that shows and modifies LEC configuration. Allows the user to set the NSAP address of the LAN Emulation Configuration Server (LECS), display the list of Emulated LANs (ELANs) configured in the LECS for this host, display the list of ELANs locally configured along with the membership state of each, and locally administer ELAN membership.

EPROM - Erasable Programmable Read Only Memory (See PROM.)

ES (End System) - a system in which an ATM connection is terminated or initiated. An originating end system initiates the ATM connection, and a terminating end system terminates the ATM connection. OAM cells may be generated and received.

Ethernet - a 10-Mbps, coaxial standard for LANs in which all nodes connect to the cable where they contend for access.

Fairness - as related to Generic Flow Control (GFC), fairness is defined as meeting all of the agreed quality of service (QoS) requirements by controlling the order of service for all active connections.

FCC - a board of commissioners appointed by the President under the Communications Act of 1934, with the authority to regulate all interstate telecommunications originating in the United States, including transmission over phone lines.

FDDI (Fiber Distributed Data Interface) - high-speed data network that uses fiber-optic as the physical medium. Operates in similar manner to Ethernet or Token Ring, only faster.

FDM (Frequency Division Multiplexing) - a method of dividing an available frequency range into parts with each having enough bandwidth to carry one channel.

FEBE (Far End Block Error) - an error detected by extracting the 4-bit FEBE field from the path status byte (G1). The legal range for the 4-bit field is between 0000 and 1000, representing zero to eight errors. Any other value is interpreted as zero errors.

FERF (Far End Receive Failure) - a line error asserted when a 110 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line FERF is removed when any pattern other than 110 is detected in these bits for five consecutive frames.

FIFO (First-In, First-Out) - a method of coordinating the sequential flow of data through a buffer.

Flag - a specific bit pattern used to identify the beginning or end of a frame.

Frame - a variable length group of data bits with a specific format containing flags at the beginning and end to provide demarcation.

Frame Relay - a fast packet switching protocol based on the LAPD protocol of ISDN that performs routing and transfer with less overhead processing than X.25.

FT-PNNI (*ForeThought* PNNI) - a FORE Systems routing and signalling protocol that uses private ATM (NSAP) addresses; a precursor to ATM Forum PNNI (see PNNI).

ftp (File Transfer Protocol) - a TCP/IP protocol that lets a user on one computer access, and transfer data to and from, another computer over a network. ftp is usually the name of the program the user invokes to accomplish this task.

GCRA (Generic Cell Rate Algorithm) - an algorithm which is employed in traffic policing and is part of the user/network service contract. The GCRA is a scheduling algorithm which ensures that cells are marked as *conforming* when they arrive when expected or later than expected and *non-conforming* when they arrive sooner than expected.

GFC (Generic Flow Control) - the first four bits of the first byte in an ATM cell header. Used to control the flow of traffic across the User-to-Network Interface (UNI), and thus into the network. Exact mechanisms for flow control are still under investigation and no explicit definition for this field exists at this time. (This field is used only at the UNI; for NNI-NNI use (between network nodes), these four bits provide additional network address capacity, and are appended to the VPI field.)

GIO - a proprietary bus architecture used in certain Silicon Graphics, Inc. workstations.

Header - protocol control information located at the beginning of a protocol data unit.

HEC (Header Error Control) - a CRC code located in the last byte of an ATM cell header that is used for checking cell integrity only.

HIPPI (High Performance Parallel Interface) - an ANSI standard that extends the computer bus over fairly short distances at speeds of 800 and 1600 Mbps.

HPUX - the Hewlett-Packard version of UNIX.

HSSI (High-Speed Serial Interface) - a serial communications connection that operates at speeds of up to 1.544 Mbps.

Hub - a device that connects to several other devices, usually in a star topology.

ICMP (Internet Control Message Protocol) - the protocol that handles errors and control messages at the IP layer. ICMP is actually a part of the IP protocol layer. It can generate error messages, test packets, and informational messages related to IP.

IEEE (Institute of Electrical and Electronics Engineers) - the world's largest technical professional society. Based in the U.S.A., the IEEE sponsors technical conferences, symposia & local meetings worldwide, publishes nearly 25% of the world's technical papers in electrical, electronics & computer engineering, provides educational programs for its members, and promotes standardization.

IETF (Internet Engineering Task Force) - a large, open, international community of network designers, operators, vendors and researchers whose purpose is to coordinate the operation, management and evolution of the Internet to resolve short- and mid-range protocol and architectural issues.

ILMI (Interim Local Management Interface) - the standard that specifies the use of the Simple Network Management Protocol (SNMP) and an ATM management information base (MIB) to provide network status and configuration information.

Interface Data - the unit of information transferred to/from the upper layer in a single interaction across a SAP. Each Interface Data Unit (IDU) controls interface information and may also contain the whole or part of the SDU.

internet - while an internet is a network, the term "internet" is usually used to refer to a collection of networks interconnected with routers.

Internet - (note the capital "I") the largest internet in the world including large national backbone nets and many regional and local networks worldwide. The Internet uses the TCP/IP suite. Networks with only e-mail connectivity are not considered on the Internet.

Internet Addresses - the numbers used to identify hosts on an internet network. Internet host numbers are divided into two parts; the first is the network number and the second, or local, part is a host number on that particular network. There are also three classes of networks in the Internet, based on the number of hosts on a given network. Large networks are classified as Class A, having addresses in the range 1-126 and having a maximum of 16,387,064 hosts. Medium networks are classified as Class B, with addresses in the range 128-191 and with a maximum of 64,516 hosts. Small networks are classified as Class C, having addresses in the range 192-254 with a maximum of 254 hosts. Addresses are given as dotted decimal numbers in the following format:

nnn.nnn.nnn.nnn

In a Class A network, the first of the numbers is the network number, the last three numbers are the local host address.

In a Class B network, the first two numbers are the network, the last two are the local host address.

In a Class C network, the first three numbers are the network address, the last number is the local host address.

The following table summarizes the classes and sizes:

Class	First #	Max# Hosts
A	1-126	16,387,064
B	129-191	64,516
C	192-223	254

Network mask values are used to identify the network portion and the host portion of the address. For:

Class A - the default mask is 255.0.0.0

Class B - the default mask is 255.255.0.0

Class C - the default mask is 255.255.255.0

Subnet masking is used when a portion of the host ID is used to identify a subnetwork. For example, if a portion of a Class B network address is used for a subnetwork, the mask could be set as 255.255.255.0. This would allow the third byte to be used as a subnetwork address. All hosts on the network would still use the IP address to get on the Internet.

IP (Internet Protocol) - a connectionless, best-effort packet switching protocol that offers a common layer over dissimilar networks.

IPX Protocol (Internetwork Packet Exchange) - a NetWare protocol similar to the Xerox Network Systems (XNS) protocol that provides datagram delivery of messages.

IS (Intermediate system) - a system that provides forwarding functions or relaying functions or both for a specific ATM connection. OAM cells may be generated and received.

ISA Bus - a bus standard developed by IBM for expansion cards in the first IBM PC. The original bus supported a data path only 8 bits wide. IBM subsequently developed a 16-bit version for its AT class computers. The 16-bit AT ISA bus supports both 8- and 16-bit cards. The 8-bit bus is commonly called the PC/XT bus, and the 16-bit bus is called the AT bus.

ISDN (Integrated Services Digital Network) - an emerging technology that is beginning to be offered by the telephone carriers of the world. ISDN combines voice and digital network services into a single medium or wire.

ISO (International Standards Organization) - a voluntary, non treaty organization founded in 1946 that is responsible for creating international standards in many areas, including computers and communications.

Isochronous - signals carrying embedded timing information or signals that are dependent on uniform timing; usually associated with voice and/or video transmission.

Jumper - a patch cable or wire used to establish a circuit, often temporarily, for testing or diagnostics; also, the devices, shorting blocks, used to connect adjacent exposed pins on a printed circuit board that control the functionality of the card.

LAN (Local Area Network) - a data network intended to serve an area of only a few square kilometers or less. Because the network is known to cover only a small area, optimizations can be made in the network signal protocols that permit higher data rates.

lane - a program that provides control over the execution of the LAN Emulation Server (LES), Broadcast/Unknown Server (BUS), and LAN Emulation Configuration Server (LECS) on the local host.

LAN Access Concentrator - a LAN access device that allows a shared transmission medium to accommodate more data sources than there are channels currently available within the transmission medium.

Layer Entity - an active layer within an element.

Layer Function - a part of the activity of the layer entities.

Layer Service - a capability of a layer and the layers beneath it that is provided to the upper layer entities at the boundary between that layer and the next higher layer.

Layer User Data - the information transferred between corresponding entities on behalf of the upper layer or layer management entities for which they are providing services.

le - a FORE program that implements both the LAN Emulation Server (LES) and the Broadcast/Unknown Server (BUS).

LEC (LAN Emulation Client) - the component in an end system that performs data forwarding, address resolution, and other control functions when communicating with other components within an ELAN.

lecs - a FORE program that implements the assignment of individual LECs to different emulated LANs.

LECS (LAN Emulation Configuration Server) - the LECS is responsible for the initial configuration of LECs. It provides information about available ELANs that a LEC may join, together with the addresses of the LES and BUS associated with each ELAN.

leq - a FORE program that provides information about an ELAN. This information is obtained from the LES, and includes MAC addresses registered on the ELAN together with their corresponding ATM addresses.

LES (LAN Emulation Server) - the LES implements the control coordination function for an ELAN. The LES provides the service of registering and resolving MAC addresses to ATM addresses.

LLC (Logical Link Control) - a protocol developed by the IEEE 802 committee for data-link-layer transmission control; the upper sublayer of the IEEE Layer 2 (OSI) protocol that complements the MAC protocol; IEEE standard 802.2; includes end-system addressing and error checking.

LOF (Loss Of Frame) - a type of transmission error that may occur in wide-area carrier lines.

looptest - a program that tests the interface for basic cell reception and transmission functionality. It is usually used for diagnostic purposes to determine if an interface is functioning properly.

LOP (Loss Of Pointer) - a type of transmission error that may occur in wide-area carrier lines.

LOS (Loss Of Signal) - a type of transmission error that may occur in wide-area carrier lines.

MAC (Media Access Control) - a media-specific access control protocol within IEEE 802 specifications; currently includes variations for Token Ring, token bus, and CSMA/CD; the lower sublayer of the IEEE's link layer (OSI), which complements the Logical Link Control (LLC).

Metasignalling - an ATM Layer Management (LM) process that manages different types of signalling and possibly semipermanent virtual channels (VCs), including the assignment, removal, and checking of VCs.

Metasignalling VCs - the standardized VCs that convey metasignalling information across a User-to-Network Interface (UNI).

MIB (Management Information Base) - the set of parameters an SNMP management station can query or set in the SNMP agent of a networked device (e.g., router).

MIC (Media Interface Connector) - the optical fiber connector that joins the fiber to the FDDI controller.

MicroChannel - a proprietary 16- or 32-bit bus developed by IBM for its PS/2 computers' internal expansion cards; also offered by others.

MTU (Maximum Transmission Unit) - the largest unit of data that can be sent over a type of physical medium.

Multi-homed - a device that has both an ATM and another network connection, typically Ethernet.

Multiplexing - a function within a layer that interleaves the information from multiple connections into one connection. (See demultiplexing.)

Multipoint Access - user access in which more than one terminal equipment (TE) is supported by a single network termination.

Multipoint-to-Point Connection - a Point-to-Multipoint Connection may have zero bandwidth from the Root Node to the Leaf Nodes, and non-zero return bandwidth from the Leaf Nodes to the Root Node. Such a connection is also known as a Multipoint-to-Point Connection.

Multipoint-to-Multipoint Connection - a collection of associated ATM VC or VP links, and their associated endpoint nodes, with the following properties:

1. All N nodes in the connection, called Endpoints, serve as a Root Node in a Point-to-Multipoint connection to all of the (N-1) remaining endpoints.

2. Each of the endpoints on the connection can send information directly to any other endpoint, but the receiving endpoint cannot distinguish which of the endpoints is sending information without additional (e.g., higher layer) information.

Network Module - ATM port interface cards which may be individually added or removed from any Cabletron ATM switch to provide a diverse choice of connection alternatives. Each network module provides between one and six full-duplex ATM physical connections to the Cabletron switch.

NMS (Network Management Station) - the system responsible for managing a network or a portion of a network. The NMS talks to network management agents, which reside in the managed nodes.

NNI (Network-to-Network Interface or Network Node Interface) - the interface between two public network pieces of equipment.

NuBus - a high-speed bus used in the Macintosh family of computers, structured so that users can put a card into any slot on the board without creating conflict over the priority between those cards

OAM (Operation and Maintenance) Cell - a cell that contains ATM LM information. It does not form part of the upper layer information transfer.

OpenView - Hewlett-Packard's network management software.

OSI (Open Systems Interconnection) - the 7-layer suite of protocols designed by ISO committees to be the international standard computer network architecture.

OSPF (Open Shortest Path First) Protocol - a routing algorithm for IP that incorporates least-cost, equal-cost, and load balancing.

Out-of-Band Management - refers to switch configuration via the serial port or over Ethernet, not ATM.

Packet Switching - a communications paradigm in which packets (messages) are individually routed between hosts with no previously established communications path.

PBX (Private Branch Exchange) - a private phone system (switch) that connects to the public telephone network and offers in-house connectivity. To reach an outside line, the user must dial a digit like 8 or 9.

PCI (Peripheral Component Interconnect) - a local-bus standard created by Intel.

PCM (Pulse Code Modulation) - a modulation scheme that samples the information signals and transmits a series of coded pulses to represent the data.

PDN (Public Data Network) - a network designed primarily for data transmission and intended for sharing by many users from many organizations.

PDU (Protocol Data Unit) - a unit of data specified in a layer protocol and consisting of protocol control information and layer user data.

Peak Cell Rate - at the PHY Layer SAP of a point-to-point VCC, the Peak Cell Rate $R_p$is the inverse of the minimum inter-arrival time T_0 of the request to send an ATM-SDU.

Peer Entities - entities within the same layer.

PHY (Physical Layer) - the actual cards, wires, and/or fiber-optic cabling used to connect computers, routers, and switches.

Physical Layer (PHY) Connection - an association established by the PHY between two or more ATM-entities. A PHY connection consists of the concatenation of PHY links in order to provide an end-to-end transfer capability to PHY SAPs.

PMD (Physical Medium Dependent) - a sublayer concerned with the bit transfer between two network nodes. It deals with wave shapes, timing recovery, line coding, and electro-optic conversions for fiber based links.

PNNI (Private Network Node Interface or Private Network-to-Network Interface) - a protocol that defines the interaction of private ATM switches or groups of private ATM switches

ping (Packet Internet Groper) - a program used to test reachability of destinations by sending them an ICMP echo request and waiting for a reply.

Point-to-Multipoint Connection - a collection of associated ATM VC or VP links, with associated endpoint nodes, with the following properties:

1. One ATM link, called the Root Link, serves as the root in a simple tree topology. When the Root node sends information, all of the remaining nodes on the connection, called Leaf nodes, receive copies of the information.

2. Each of the Leaf Nodes on the connection can send information directly to the Root Node. The Root Node cannot distinguish which Leaf is sending information without additional (higher layer) information. (See the following note for Phase 1.)

3. The Leaf Nodes cannot communicate directly to each other with this connection type.

Note: Phase 1 signalling does not support traffic sent from a Leaf to the Root.

Point-to-Point Connection - a connection with only two endpoints.

Primitive - an abstract, implementation-independent interaction between a layer service user and a layer service provider.

PROM (Programmable Read-Only Memory) - a chip-based information storage area that can be recorded by an operator but erased only through a physical process.

Protocol - a set of rules and formats (semantic and syntactic) that determines the communication behavior of layer entities in the performance of the layer functions.

Protocol Control Information - the information exchanged between corresponding entities using a lower layer connection to coordinate their joint operation.

Proxy - the process in which one system acts for another system to answer protocol requests.

Proxy Agent - an agent that queries on behalf of the manager, used to monitor objects that are not directly manageable.

PSN (Packet Switched Network) - a network designed to carry data in the form of packets. The packet and its format is internal to that network.

PT (Payload Type) - bits 2...4 in the fourth byte of an ATM cell header. The PT indicates the type of information carried by the cell. At this time, values 0...3 are used to identify various types of user data, values 4 and 5 indicate management information, and values 6 and 7 are reserved for future use.

PVC (Permanent Virtual Circuit (or Channel)) - a circuit or channel through an ATM network provisioned by a carrier between two endpoints; used for dedicated long-term information transport between locations.

Q.2931 - Derived from Q.93B, the narrowband ISDN signalling protocol, an ITU standard describing the signalling protocol to be used by switched virtual circuits on ATM LANs.

Relaying - a function of a layer by means of which a layer entity receives data from a corresponding entity and transmits it to another corresponding entity.

RFCs (Requests For Comment) - IETF documents suggesting protocols and policies of the Internet, inviting comments as to the quality and validity of those policies. These comments are collected and analyzed by the IETF in order to finalize Internet standards.

RFI (Radio Frequency Interference) - the unintentional transmission of radio signals. Computer equipment and wiring can both generate and receive RFI.

RIP (Routing Information Protocol) - a distance vector-based protocol that provides a measure of distance, or hops, from a transmitting workstation to a receiving workstation.

RISC (Reduced Instruction Set Computer) - a generic name for CPUs that use a simpler instruction set than more traditional designs.

Router - a device that forwards traffic between networks or subnetworks based on network layer information.

SBus - hardware interface for add-in boards in later-version Sun 3 workstations.

SAP (Service Access Point) - the point at which an entity of a layer provides services to its LM entity or to an entity of the next higher layer.

SAR (Segmentation And Reassembly) - the SAR accepts PDUs from the CS and divides them into very small segments (44 bytes long). If the CS-PDU is less than 44 bytes, it is padded to 44 with zeroes. A two-byte header and

trailer are added to this basic segment. The header identifies the message type (beginning, end, continuation, or single) and contains sequence numbering and message identification. The trailer gives the SAR-PDU payload length, exclusive of pad, and contains a CRC check to ensure the SAR-PDU integrity. The result is a 48-byte PDU that fits into the payload field of an ATM cell.

SCSI (Small Computer Systems Interface) - a standard for a controller bus that connects disk drives and other devices to their controllers on a computer bus. It is typically used in small systems.

SDLC (Synchronous Data Link Control) - IBM's data link protocol used in SNA networks.

SDU (Service Data Unit) - a unit of interface information whose identity is preserved from one end of a layer connection to the other.

SEAL (Simple and Efficient Adaptation Layer) - also called AAL 5, this ATM adaptation layer assumes that higher layer processes will provide error recovery, thereby simplifying the SAR portion of the adaptation layer. Using this AAL type packs all 48 bytes of an ATM cell information field with data. It also assumes that only one message is crossing the UNI at a time. That is, multiple end-users at one location cannot interleave messages on the same VC, but must queue them for sequential transmission.

Segment - a single ATM link or group of interconnected ATM links of an ATM connection.

Semipermanent Connection - a connection established via a service order or via network management.

SGMP (Simple Gateway Management Protocol) - the predecessor to SNMP.

Shaping Descriptor - *n* ordered pairs of GCRA parameters (I,L) used to define the negotiated traffic shape of an APP connection. The traffic shape refers to the load-balancing of a network. In this context, load-balancing means configuring the data flows to maximize the efficiency of the network.

SIR (Sustained Information Rate) - the long-term average data transmission rate across the User-to-Network Interface.

SMDS (Switched Multimegabit Data Service) - a high-speed, datagram-based, public data network service expected to be widely used by telephone companies in their data networks.

SMTP (Simple Mail Transfer Protocol) - the Internet electronic mail protocol used to transfer electronic mail between hosts.

SNAP - SubNetwork Access Protocol

SNMP (Simple Network Management Protocol) - the Internet standard protocol for managing nodes on an IP network.

snmpd - an SMNP agent for a given adapter card.

SONET (Synchronous Optical Network) - a new and growing body of standards that defines all aspects of transporting and managing digital traffic over optical facilities in the public network.

Source Traffic Descriptor - a set of traffic parameters belonging to the ATM Traffic Descriptor used during the connection set-up to capture the intrinsic traffic characteristics of the connection requested by the source.

Spanning Tree Protocol - provides loop-free topology in a network environment where there are redundant paths.

SPANS (Simple Protocol for ATM Network Signalling) - FORE Systems' proprietary signalling protocol used for establishing SVCs between FORE Systems equipment.

SPARC (Scalable Processor Architecture Reduced instruction set Computer) - a powerful workstation similar to a reduced-instruction-set-computing (RISC) workstation.

SPE (Synchronous Payload Envelope) - the payload field plus a little overhead of a basic SONET signal.

SPVC (Smart PVC) - a generic term for any communications medium which is permanently provisioned at the end points, but switched in the middle. In ATM, there are two kinds of SPVCs: smart permanent virtual path connections (SPVPCs) and smart permanent virtual channel connections (SPVCCs).

Static Route - a route that is entered manually into the routing table.

STM (Synchronous Transfer Mode) - a transport and switching method that depends on information occurring in regular and fixed patterns with respect to a reference such as a frame pattern.

STP (Shielded Twisted Pair) - two or more insulated wires that are twisted together and then wrapped in a cable with metallic braid or foil to prevent interference and offer noise-free transmissions.

STS (Synchronous Transport Signal) - a SONET electrical signal rate.

Sublayer - a logical subdivision of a layer.

SVC (Switched Virtual Circuit (or Channel)) - a channel established on demand by network signalling, used for information transport between two locations and lasting only for the duration of the transfer; the datacom equivalent of a dialed telephone call.

Switched Connection - a connection established via signalling.

Symmetric Connection - a connection with the same bandwidth value specified for both directions.

Synchronous - signals that are sourced from the same timing reference and hence are identical in frequency.

Systems Network Architecture (SNA) - a proprietary networking architec

TAXI - Transparent Asynchronous Transmitter/Receiver Interface

TC (Transmission Convergence) - generates and receives transmission frames and is responsible for all overhead associated with the transmission frame. The TC sublayer packages cells into the transmission frame.

TCP (Transmission Control Protocol) - a specification for software that bundles and unbundles sent and received data into packets, manages the transmission of packets on a network, and checks for errors.

TCP/IP (Transmission Control Protocol/Internet Protocol) - a set of communications protocols that has evolved since the late 1970s, when it was first developed by the Department of Defense. Because programs supporting these protocols are available on so many different computer systems, they have become an excellent way to connect different types of computers over networks.

TDM (Time Division Multiplexing) - a traditional digital multiplexing in which a signal occupies a fixed, repetitive time slot within a higher-rate signal.

Token Ring - a network access method in which the stations circulate a token. Stations with data to send must have the token to transmit their data.

Traffic - the calls being sent and received over a communications network. Also, the packets that are sent on a data network.

Trailer - the protocol control information located at the end of a PDU.

Transit Delay - the time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.

trap - a program interrupt mechanism that automatically updates the state of the network to remote network management hosts. The SNMP agent on the switch supports these SNMP traps.

UBR (Unspecified Bit Rate) - a type of traffic that is not considered time-critical (e.g., ARP messages, pure data), allocated whatever bandwidth is available at any given time. UBR traffic is given a "best effort" priority in an ATM network with no guarantee of successful transmission.

UDP (User Datagram Protocol) - the TCP/IP transaction protocol used for applications such as remote network management and name-service access; this lets users assign a name, such as "RVAX*2,S," to a physical or numbered address.

Unassigned Cells - a cell identified by a standardized virtual path identifier (VPI) and virtual channel identifier (VCI) value, which has been generated and does not carry information from an application using the ATM Layer service.

UNI (User-to-Network Interface) - the physical and electrical demarcation point between the user and the public network service provider.

UNI 3.0 - the User-to-Network Interface standard set forth by the ATM Forum that defines how private customer premise equipment interacts with private ATM switches.

UTP (Unshielded Twisted Pair) - a cable that consists of two or more insulated conductors in which each pair of conductors are twisted around each other. There is no external protection and noise resistance comes solely from the twists.

VBR (Variable Bit Rate) - a type of traffic that, when sent over a network, is tolerant of delays and changes in the amount of bandwidth it is allocated (e.g., data applications).

VC (Virtual Channel (or Circuit)) - a communications path between two nodes identified by label rather than fixed physical path.

VCC (Virtual Channel Connection) - a unidirectional concatenation of VCLs that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer (i.e., a higher layer or ATMM-entity) for processing signify the endpoints of a VCC.

VCI (Virtual Channel Identifier) - the address or label of a VC.

VCL (Virtual Channel Link) - a means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.

VINES (Virtual Network Software) - Banyan's network operating system based on UNIX and its protocols.

Virtual Channel Switch - a network element that connects VCLs. It terminates VPCs and translates VCI values. The Virtual Channel Switch is directed by Control Plane functions and relays the cells of a VC.

Virtual Path Switch - a network element that connects VPLs, it translates VPI (not VCI) values and is directed by Control Plane functions. The Virtual Path Switch relays the cells of a Virtual Path.

VPT (Virtual Path Terminator) - a system that unbundles the VCs of a VP for independent processing of each VC.

VP (Virtual Path) - a unidirectional logical association or bundle of VCs.

VPC (Virtual Path Connection) - a concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.

VPDN (Virtual Private Data Network) - a private data communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

VPI (Virtual Path Identifier) - the address or label of a particular VP.

VPL (Virtual Path Link) - a means of unidirectional transport of ATM cells between the point where a VPI value is assigned and the point where that value is translated or removed.

VPN (Virtual Private Network) - a private voice communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

VT (Virtual Tributary) - a structure used to carry payloads such as DS1s that run at significantly lower rates than STS-1s.

WAN (Wide-Area Network) - a network that covers a large geographic area.

X.25 - a well-established data switching and transport method that relies on a significant amount of processing to ensure reliable transport over metallic media.